

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 869 443 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.10.1998 Bulletin 1998/41

(51) Int. Cl.⁶: G06F 17/30

(21) Application number: 97304294.8

(22) Date of filing: 19.06.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE
Designated Extension States:
AL LT LV RO SI

(30) Priority: 04.04.1997 US 826560

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)

(72) Inventors:
• Shaw, George H.J.
Woodinville, WA 98072 (US)

• O'Rourke, Thomas J.
90800 Oulu (FI)
• Woodruff, Bryan A.
New Bend, WA 98045 (US)

(74) Representative:
Belcher, Simon James et al
Urquhart-Dykes & Lord
Tower House
Merrion Way
Leeds LS2 8PA (GB)

(54) **Method and computer program product for synchronizing the processing of multiple data streams and matching disparate processing rates using a standardized clock mechanism**

(57) A method and computer program product for synchronizing processing between two or more data streams (e.g., video and sound input) and for rate matching between two different hardware clocks that may drift with respect to one another (e.g., an originating clock represented in a timestamped data stream versus a clock actually rendering the data) in a system of interconnected software drivers running in kernel mode. The present invention overcomes the coordination complexity and inaccuracies in the prior art by providing a clocking mechanism in a system wherein multiple drivers having input and output connection pin instances are chained together. The clocking mechanism synchronizes between data streams by providing a master clock on an input pin instance of a driver that is used to synchronize with other input pin instances on other drivers and "slave" clocks. Synchronization is achieved through event notification or stream position queries so that corresponding frames of data in separate streams are rendered together (e.g., video frames with corresponding sound track). Rate matching is achieved through monitoring a physical clock progression in comparison with a series of data stream timestamps thereby allowing adjustments to match the different clock rates. A common physical clock (e.g., PC clock) can be used as a reference for a component to translate a particular clock time to a time shared by all components with a minimum of error.

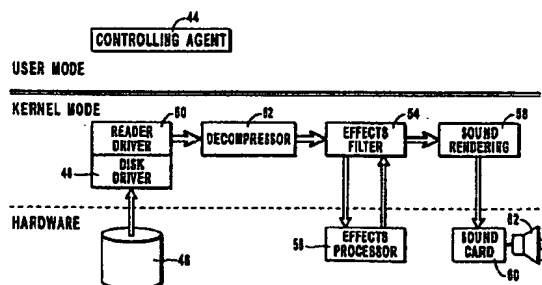


FIG. 2

EP 0 869 443 A2

Description

The field of the present invention is methods and computer program products relating to timing issues in multimedia applications, such as stream synchronization and rate matching between a generating clock and a rendering clock. The present invention also relates to standardized timing mechanisms in software drivers. More specifically, the present invention is directed to methods and computer program products for providing timing and clock mechanisms used by multiple entities, such as interconnected software drivers, in a standardized fashion so that multiple streams of processed data may be synchronized and rates between different hardware clocks may be matched.

When processing multimedia data, such as digitized sound and video data, it is common to process a continuous stream of media samples. The data stream also has time interval information associated with the data either by convention or by timestamp information. Timestamp information will tell a processing component when a particular sample is to be processed or can be used as a tracking reference to assure processing is progressing at the rate specified between timestamps.

Data organization, format, and convention can also convey time interval information used during processing. For example, a file of video data in a 30 frames per second (fps) format and convention will be understood by a processing component, having knowledge of the convention and format, that each sample is processed at 1/30 of a second intervals.

One problem encountered in processing media streams is that of synchronizing two or more streams so that they are rendered together. For example, an audio stream representing a soundtrack and an accompanying stream of video data need to be synchronized in order to produce a coherent multimedia presentation.

One method of synchronization utilizes a clock mechanism having a value that is based on a physical or hardware clock such as the PC clock. A data stream processor may query the value of this clock mechanism in order to synchronize processing or rendering of the stream. While synchronization occurs with respect to the hardware clock, it does not necessarily synchronize two different streams due to processing events and delays that may occur independently on each stream.

Another method of synchronization utilizes a clock mechanism that provides a time value that is based on presentation timestamp values on a "master" data stream while the stream is being processed. When the master stream is paused or otherwise stopped, the time value alternates to be based off a physical or hardware clock such as the PC clock. While the processor of a second data or media stream may query the value of this clock mechanism, further facility is provided so that the clock mechanism will notify the processor at a specific time on the physical clock or in certain intervals (e.g., every 10 ms), again based on the physical clock.

Since event notification is based on the physical clock, such events may continue to occur even when processing is paused or stopped on the master data or media stream. Notifications may be temporarily disabled but at the cost of extra processing overhead and increased processing code complexity. Even when notifications are specifically disabled, timing circumstances may be such that a notification still is erroneously made.

Another problem associated with processing media or data streams is that of matching the rates of different hardware clocks associated with the media stream. One hardware clock will be used in the generation of the media stream media samples and making the timestamp information associated with each sample. A completely different hardware clock, however, will be used to render the media stream. Though the two hardware clocks may ostensibly run at the same rate, variations in physical clock rates may impact processing.

Rate matching, then, is the concept of compensating for actual frequency variations in physical clocks. For example, if a stream of live audio is being received through a network connection for rendering on local audio hardware, the rate at which the stream is being received may not match the rate at which the local audio hardware actually renders the stream. This is because the crystal on the audio hardware has nothing to do with the crystal used on the audio hardware on the other end of the network connection.

In some instances it is possible to perform fine adjustments on the local audio hardware in order to attempt to move it closer in frequency to what the source is producing. Not all audio hardware will allow this, however, and those that do have certain frequency resolution constraints that limit the accuracy of an adjustment and may have limits as to the amount of adjustment allowed.

Processing "drift" occurs due to this rate differential and ways are sought to determine the amount of drift accurately and quickly. Further, compensating for the drift optimally occurs in frequent minute increments, as opposed to infrequent gross adjustments, in order to avoid noticeable perception by an end user. Such perceptible adjustments come in the form of "skipping" of sound data and/or "jerkiness" of video images.

Yet another problem encountered with processing media or data streams is that of translating one time frame to another. This often occurs during synchronization as the presentation times of a "slave" stream are converted to the time frame of the master stream. Ways are sought to achieve such translations with minimal error introduction.

In a multimedia environment, it is advantageous to interconnect software drivers so that processing may occur in software drivers that run in an operating system mode with a much higher run priority and little security protection to allow

access to actual hardware that the drivers, in many instances, directly manipulate. Many applications are benefited by running in this looser and more performance-oriented mode, generally referred throughout, in Windows NT terminology, as "kernel mode." Other robust operating systems will have a functionally equivalent mode. Such interconnected software drivers have the same need for synchronizing multiple media streams and for rate matching of physical clocks.

One prime example of a program currently incapable of easily using kernel mode drivers, used throughout this application, comprises graph building functionality that allows a user to select and connect together different processing blocks, called filters, to successively manipulate a stream of multimedia data. The data typically is a series of samples representing sound or video and the processing blocks may include decompression processing for compressed data, special effects processing, CODEC functionality, rendering the data into analog signals, etc.

Such filters are typically located in user mode so that the graph builder portion of the program may interconnect and control their operation and be responsive to user input and rearrangement of processing blocks. Because of the consistent stream nature of multimedia data and the generation of large quantities of data, performance is a critical issue. In a general purpose operating system, system performance caused by repeatedly passing/switching back and forth between user mode and kernel mode can be so degraded as to prohibit certain multimedia applications.

Ways to standardize the interconnection of kernel mode filters would be desirable and would further benefit by a standardized clock mechanism to assist in solving the synchronization problem associated with processing multiple media streams simultaneously. Moreover, such standardized clock mechanism should also address and solve the rate matching problem associated with originating the data and rendering the data using different physical clocks.

The present invention seeks to synchronize processing of multiple data streams, each composed of data samples having time interval information associated therewith.

It seeks to facilitate matching the processing rate of a data stream where the data stream was created by an originating physical clock as represented in the stream data samples and a rendering physical clock associated with a piece of hardware.

It seeks to provide a useful clock mechanism having a standardized interface for use in a system of interconnected kernel mode software drivers.

The invention seeks to facilitate translation from time on one clock mechanism to time on another clock mechanism.

Accordingly, the invention provides a method and a computer program product for synchronizing the processing of multiple data streams and matching disparate processing rates using a standardized clock mechanism is provided. The clock mechanism can be implemented in a variety of different ways and is explained throughout in the context of interconnectable kernel mode drivers providing efficient data stream processing.

The clock mechanism of the present invention maintains three different time values and makes them available to other components. The time values may be queried or the clock mechanism may send notifications based on arriving at a given time value or at periodic intervals.

The first time value available at the clock mechanism of the present invention is a positional time value that is based on the time interval information associated with a data stream and reflects the "position" within the data stream being processed. The second is a physical time value based on an actual hardware oscillator or clock, such as the PC clock or a rendering processor clock. Finally, the third is a correlated time value that provides a designated time value, such as the positional time value, together with a reference time value based on a commonly available hardware clock such as the PC clock.

With the clock mechanism as a master clock, the positional time value can be used by other processing components processing other data streams for synchronization. Since the positional time value is based on processing of a data stream, when the processing of the underlying data stream stops the clock does as well. Furthermore, timing event notifications will also be frozen on an absolute time basis, but will continue to be perfectly timed with respect to the positional time value. This provides a great benefit in that "paused" time need not be tracked, manipulated, or otherwise be of concern. "Slave" clock mechanisms and components will base processing on time as it is provided by the positional time value.

The physical time value of the underlying hardware clock associated with a data stream renderer or other processor can be used to match the rate of the data stream being processed with the processing rate of the rendering processor. Since the data stream has time interval information associated therewith, the originating processor's processing rate is inherent therein and may be ascertained and compared to the rendering processors actual processing rate as found in the physical time value. This allows any processing rate "drift" to be quickly and accurately determined thereby allowing processing adjustments to be made with less noticeable effect on the rendered data.

Finally, the correlated time value provides a means for other clock mechanisms and processing components to translate time information from one basis to another. The common hardware clock must be available to other clock mechanisms or components so that it may be used as a common reference. Since the correlated time is returned or accessed as an atomic operation, any errors introduced are minimal and non-cumulative.

Two variations of correlated time are supported in the described embodiment, both using the PC clock as a reference. One uses media time as the designated time and the other uses the physical time as the designated time.

Throughout this application, correlated time used alone could refer to either of the two variations or even other variations that will be apparent to those skilled in the art. When distinguishing between the variations, the terms correlated media time and correlated physical time will be used with the respective designated time compared to the common reference.

One embodiment of the present invention occurs as part of connecting kernel mode drivers in a standardized fashion. A given driver or filter will support and define connection point "pin factories" that may be used to instantiate connection pin instances that are interconnected to other pin instances on other drivers to allow processing messages to be consecutively processed in kernel mode by the drivers without necessary resort to a user mode agent.

A third party agent desiring to connect compliant drivers will query the drivers of their capabilities. Such capabilities include what kinds of connection pin instances may be instantiated, including their relevant characteristics, such as type of data handled, data formats, transfer rates, medium or mode of transfer, input or output nature of a connection pin instance, etc. Also queried will be the data buffering requirements and capabilities for allocating buffers available at each connection pin factory.

Once a third party agent, typically running in user mode, has queried the capabilities of one or more compliant drivers, the agent will determine the best connection characteristics for "chaining" multiple drivers together so that data may be optimally processed between them. This determination step occurs after all driver capabilities have been queried so that the optimal connection criteria may be selected.

For each connection pin instance, the third party agent will also determine whether a buffer allocator need be created on a particular pin instance. Again, this is done after having queried all the different filters and pin factories prior to making any interconnections.

The third party agent then interconnects the drivers by creating a connection pin instance based on the connection pin factories on each driver. The agent will specify a data format and a connection format as part of the connection pin instance creation. In an exemplary embodiment implemented under the NT operating system, an actual connection pin instance is created by a create I/O operation that returns a handle to a "file". The create I/O request will contain the driver instance handle and reference to a data structure indicating data format and connection format for the connection pin instance.

Furthermore, reference to previously created connection pin instances (*e.g.*, input pins or IRP "sink" pins) will be specified in requests for creating other connection pin instances (*e.g.*, output pins or IRP "source" pins) in order to effectuate a connection between connection pin instances. After a source pin instance is created using a reference to an input pin instance having a buffer allocator, indication is made with the source pin instance of the buffer allocator so that data may be transmitted into the new buffer from the existing buffer. If no reference is indicated, the source pin instance will leave the data in the existing data after processing.

In order to create a compliant driver, a driver developer will support certain standard facilities to allow a user mode agent to query capabilities and make interconnections between drivers. In one embodiment, built on the Windows NT operating system, this is achieved by use of "sets" (*i.e.*, property, method, and event sets) that implement the desired functionality.

A set is logically defined as having a GUID (globally unique identifier) to identify the set as a whole and a RUID (relatively unique identifier, *e.g.*, relative within the set itself) for each element of functionality within the set. Each set is associated with only one or two IOCTLs (IO Controls), and an IOCTL combined with a set specification controls all interaction with the driver.

As currently embodied, three types of sets are utilized, namely, property sets, method sets, and event sets. Property sets are used for managing values or settings within the driver, such as sound volume, transfer rate, etc., and use a single IOCTL with a flag indicating whether the call is getting a property value and or setting a property value. Method sets are used for managing the operations that a driver may perform, such as allocating memory, flushing buffers, etc., and uses a single IOCTL to call the specified method. Event sets are used for managing events associated with driver processing, such as device change notification, data starvation notification, etc., and uses two IOCTLs, one for enabling a specified event and one for disabling a specified event.

To use a set, an I/O control operation is initiated using the specified IOCTL and reference to a data structure having the set GUID, RUID, and other necessary data. For example, setting a volume property on a sound card driver would entail an I/O control operation using a set property IOCTL, specifying the appropriate GUID for the property set having the volume setting, indicating the specific RUID within that set indicates the volume property, and containing the new volume setting value.

To query the sets supported, a null GUID is used along with a query flag on a specified IOCTL for a particular set type (*e.g.*, property set IOCTL, method IOCTL, or event enable IOCTL) and a list of set GUIDs supported will be returned. To query supported properties, methods, or events within a given set, the set GUID, set type IOCTL, and a query flag are used with the operation returning a list of supported RUIDs.

By using the generic set mechanism, a minimum of functionality may be implemented to support a compliant driver but still allow unlimited extensibility. A set may be defined in a written specification that can be independently coded by

a multitude of different driver developers to create a system of interoperable and interconnectable drivers as long as particular sets are implemented. Furthermore, the specification can define mandatory properties, methods, and events that must be supported as well as optional properties, methods, and events that can be implemented depending on the driver functions and advanced capabilities. In addition to the basic minimum commonality required, driver developers may incorporate additional functionality by defining their own sets and assigning them a GUID. By being able to enumerate supported functionality (*i.e.*, make queries for supported GUIDs and RUIDs), a caller, such as a third party controlling agent, can adjust expectations or make appropriate compensation depending on the capabilities of the underlying filters.

An embodiment of a clock mechanism having a positional time value, a physical time value, a correlated positional time value, and a correlated physical time value may be created on a compliant connection pin instance. A "file" is opened on the particular device in the I/O indicating the handle of the file object representing the connection pin instance as parent and a GUID string value indicating the clock mechanism file type. A handle is returned to a file object representing the clock mechanism and the filter will support, through this file object, property sets for providing the respective time values. The positional time value will be based on the data stream processed by the filter and passing through the particular connection pin instance. The physical time will be based on a hardware clock or oscillator associated with the filter, if present, or the PC clock. Both of the correlated time values will typically use the PC clock as the reference value due to its widespread availability.

The clock mechanism may be used as a "master" allowing other clock mechanisms or components to "slave" off or synchronized with its time tracking. Furthermore, a default implementation based on the PC clock is provided.

As used herein, the term "user mode" refers to a level of operation in an operating system where most user written programs run. The user mode level of operation is typically the most secure level and has a significant amount of overhead to prevent one application program or process from interfering with another application program or process. Furthermore, access to system resources is highly controlled through specific interfaces and run priority is generally one of the lowest, if not the lowest.

As used herein, the term "kernel mode" refers to a level of operation in an operating system having significantly less restrictions than the user mode level of operation. Examples of kernel mode programs or processes would include software drivers for controlling hardware components. Typically, kernel mode programs are performance sensitive, and therefore, have less operational overhead than user mode programs. Furthermore, access to hardware and many system resources is unrestricted or much less restricted than for user mode programs. In many instances, program code running in kernel mode relies on programmer discipline and conformity to convention in order to establish good system behaviour (*e.g.*, not disrupting another program's address space, etc.). Another term used for kernel mode is "trusted" code.

As used herein the term "driver" refers to software driver programs typically running in kernel mode. The term driver may also refer to the actual executable program that is loaded onto the operating system or a portion thereof that imparts certain functionality. Drivers are in many instances, though not necessarily, associated with some form of hardware.

As used herein, the term "filter" refers to a portion of the functionality found within a software driver, including the entire driver itself, where connection points may be exposed for sending data through the filter. For example, a software driver may support a number of different filters or may have one single function. Furthermore, a number of filters from different drivers that are internally connected together and externally exposing connection points for input and output may collectively be referred to as a single filter. Also, in a more generic sense, the term filter may refer to the operation performed, such as decompression, etc., regardless of whether that occurs in a software driver filter running in kernel mode or another piece of program code running in user mode.

As used herein, the term "driver object" refers to an operating system entity, defined by the operating system, for managing and making known a software driver as a system resource.

As used herein, the term "device object" refers to a system level entity defined by the system, for making known a portion of a driver's functionality available for use as a system resource and defines the driver functionality and availability to other system components. Both the driver objects and device objects are typically created upon loading and initialization of the driver software.

As used herein, the term "file object" refers to an operating system entity, defined by the system, for managing an invocation of a resource specified by a device object. A file object provides a context for usage of a driver object. Furthermore, a file object may be hierarchically related to another file object if the previous file object is designated as a "parent" during the creation of the new file object. File objects are typically used in managing all I/O operations that operate on a stream of data.

As used herein, the term "data" refers to any information that is processed through the interconnected kernel mode filters. Such data includes media data representing video, audio, text, MIDI, etc. but may also include control information or parameters for other applications. For example, a kernel mode filter graph may be used in process control operations where the control information passed between the different filters is used to develop control signals for actuating

machinery. While examples are given of media processing systems, other applications could in like manner benefit from the system of interconnected kernel mode filters explained herein.

Throughout this specification, the present invention is described within the context of the Windows NT™ operating system available from Microsoft®. Furthermore, familiarity with the Windows NT I/O architecture is presumed in order to understand the preferred embodiment explained herein. A good tutorial of the I/O system as well as the NT operating system in general can be found in the book "Inside Windows NT" written by Helen Custer and published by Microsoft Press.

While the discussion of the drivers and system entities such as file objects, device objects and driver objects are explained herein within the context of how they operate in the Windows NT operating system, those skilled in the art will appreciate that the present invention may be implemented on other operating systems having analogous components, whether or not they use the same terminology. For example, should another operating system have an entity that operates as a file object, it could be interpreted as a file object regardless of its actual title.

The present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a prior art data flow diagram showing a system of interconnected filters and drivers under the direction of a controlling agent for bringing sound data from a disk file, processing the sound data in some form, and rendering the sound data to be played through a speaker.

Figure 2 shows a system according to the present invention having the same purpose as that shown in Figure 1 to read sound data from a disk drive, process that data, and render that data to be heard on a speaker, wherein the processing filters and rendering are handled by interconnected kernel mode drivers, again under the direction of a controlling agent.

Figure 3 is a vertical relationship model showing the relationships between driver objects, device objects and file objects as created and used in an operating system.

Figures 4A, 4B and 4C are logical block diagrams of a driver object, device object, and file object, respectively, showing their logical relationship with the data structures and program code to route messages to appropriate process handling code and to validate the creation of new file objects according to the system of the present invention. Figure 5 is a flowchart showing the initial set up of the routing and validation componentry and the processing of I/O messages by the kernel mode drivers.

Figure 6 is a flowchart showing in more detail the processing of a controlling agent, the routing and validation mechanisms, and specific create handler routines for creating a new file object.

Figure 7 is a logical diagram showing the horizontal relationship between connected filters utilizing the file object structures in an operating system to effectuate such a connection in a standardized fashion.

Figure 8 is a flowchart showing the processing steps taken by a controlling agent in user mode to create and connect the kernel mode filters or drivers of Figure 7 in order to effectuate a connection for processing I/O requests received from the controlling agent with processing continuing between different drivers (filters).

Figures 9A and 9B are logical overview diagrams of the kernel mode drivers and connections used to create a chain of kernel mode filters under the direction of a user mode controlling agent to implement a system for reading sound data from a hard drive, processing the data with the kernel mode filters, and rendering the data to be heard through a speaker.

Figure 10 is a flowchart showing the processing steps for creating the interconnected kernel mode drivers for the system shown in Figures 9A and 9B.

Figures 11A and 11B illustrate the functioning of a buffer allocator mechanism. Figure 11A shows a logical arrangement and processing of the allocated buffer frames as they are passed from one processing component to another.

Figure 11B illustrates a buffer allocator being represented as a file object that is a "child" of a file object representing an input pin instance in a system of interconnected kernel mode filters. Both Figures 11A and 11B illustrate the same filter graph topology.

Figure 12 shows the buffer allocation in transitions of the system illustrated in Figures 9A and 9B utilizing buffer allocators for controlling the allocation of buffer frames.

Figure 13 is a flow chart showing the processing steps for bringing data from a disk driver through a chain of interconnected kernel mode filters and rendering the data on sound processing hardware specifically showing the operation of buffer allocators and the actual data transferring between buffers for the system shown in Figure 12.

Figure 14 is a logical block diagram illustrating how two live data streams can be synchronized to a single master clock mechanism.

Figures 15A and 15B are logical block diagrams showing the live audio system of Figure 14 as implemented using the interconnected filter system explained in detail in connection with Figure 12. Figure 15A represents the live video stream rendering filters which receives a master clock synchronization signal, and Figure 15B illustrates the live audio rendering system that has a master clock mechanism for synchronizing both streams together and for

rate matching the live audio data with the actual audio rendering hardware.

Figures 16A-16C illustrate how a clock mechanism may be used to synchronize data processing of multiple data streams, rate match a stream of data with the physical capabilities of a hardware renderer, and translate one time base to another using a common time base. Figure 16A is a flow chart showing the synchronization processing steps. Figure 16B is a flow chart showing the rate matching processing steps. Figure 16C is a flow chart showing the translation processing steps.

Referring to Figure 1, an example system is presented for reading a stream of sound data from a disk drive and rendering that sound data so that it may be heard through a speaker according to the prior art model. An amount of data is stored on hard drive 20 representing sound in the form of digitized sound samples. Alternatively, the source of the sound data stream may be digitized information coming over a phone line, digitized information from network or other communication packets, or other sources known in the art. The data stream is composed of digitized samples having time interval information associated therewith either by data format and convention or by explicit timestamp information attached to each sample. A kernel mode disk driver 22 interacts with the disk drive hardware 20 and is under control of a user mode reader program component 24. A controlling agent 26 manages the different components in order to effectuate the rendering of the sound data and may include dynamic graph building capabilities so that the different software components may be dynamically allocated in order to provide custom filtering or other processing paths as designated by an end user.

The reader component 24 will interact with disk driver 22 using a standard I/O control interface of the operating system and will cause the compressed sound data to be read from the disk drive 20 into buffers allocated in user mode as part of the user mode process address space. Next, a decompressor component 28 will decompress the compressed data into its decompressed format for processing. As shown, this step is done entirely in user mode with the attendant lower priority and process behaviour safety mechanisms.

The effects filter 30 will operate on the data to provide some special effect and will have an accompanying effects filter 32 operating in kernel mode. Furthermore, an effects processor 34 may be present or the effects filter may operate entirely in software emulating the actual hardware processor. In order to access the effects filter 32 the effects component 30 will use the system I/O control mechanism to transfer the data and control to the effects filter. Again, the kernel mode/user mode boundary is crossed in order to make this transition.

The effects filter 32 will control the effects processor 34 and cause whatever special effect is necessary or desired to be made on the data. This may entail copying all the data from the effects component 30 down to the effects filter 32 and again to the effects processor 34 depending on actual system configuration. While many software effects components have a hardware processor associated therewith, others function entirely within system software running on the host processor.

After control and the data is transferred back into user mode at the completion of the processing of the effects component 30, it is then transferred to sound rendering component 36. The sound rendering component 36 will transfer the control and data to the sound rendering driver 38 which in turn controls the sound card 40 in order to render the data, as processed and filtered, as sound through speaker 42. As can be readily seen, there exists a variety of transfers between user mode and kernel mode that add inefficiencies to the rendering of the sound data. Because of the timing sensitive nature of multimedia data, such as a continuous stream of sound or video, it is advantageous to reduce these inefficiencies and transitions of control as well as the multiple copying of data between different buffers.

One embodiment of the present invention and used throughout will consist of a service provided on the Windows NT operating system architecture. This service is broken into different software components that a user of the system will access. First, a user mode API is available that will include a routine for creating connection pin instances and other file objects representing particular functionality, such as a clock mechanism or a buffer allocation mechanism. Additionally, and more importantly, there will be a complete set of routines and data structures to assist the driver developer in making drivers that are compliant with the standardized architecture. By utilizing such facilities from the system, different driver developers may create compliant drivers that will interact with one another according to the specified architecture. User mode agents communicate with compliant drivers through an environment subsystem running in user mode that will communicate with the system services of the NT executive and the I/O manager. This is the same standard I/O mechanism for all other I/O and the present implementation of the preferred embodiment will utilize existing system services as much as possible.

The architecture of the system of Figure 1 utilizing the present invention will appear as shown in Figure 2. A controlling agent 44 will query the drivers known in order to then make interconnections according to data format and connection format to effectuate the rendering entirely in kernel mode. Furthermore, the controlling agent will receive notifications of important events so that it may exercise control as necessary. Examples of such events would include end of processing, a data starvation situation, etc.

In this configuration, the sound data is read from disk drive 46 by the disk driver 48, as before. A reader driver 50 controls disk driver 48 and is "vertically" associated with disk driver 48 according to the NT layered I/O architecture as

used in conventional fashion. The terms vertically and horizontally are used to distinguish driver connections that currently occur as part of the NT layered I/O architecture (vertical) and connections according to the interconnected kernel mode drivers made dynamically by a third party controlling agent (horizontal).

Reader driver 50 is also interconnected "horizontally" to a decompressor driver 52 according to the connection methods explained hereafter and is managed by the controlling agent 44. Decompressor 52 will perform the decompression in kernel mode before passing the data and control to the effects filter 54. The effects filter will apply the special effects utilizing an effects processor 56 as necessary before passing the data and control to the sound rendering driver 58 that controls the sound card and causes the data to be rendered as sound on speaker 62. As can be noted by reference to Figure 2, keeping processing in kernel mode represents an efficiency advantage by eliminating multiple transitions between user mode and kernel mode and by reducing the amount of overhead normally associated with processing in user mode.

Referring now to Figure 3, a logical diagram showing the hierarchal nature of system objects related to interconnected software drivers compliant with one embodiment of the present invention is shown. A driver object 64 is created to represent the executable software code image as loaded in memory. The driver code image contains the entirety of the driver's functionality, and the driver object 64 includes information regarding the image, such as its location on the system, the kinds of devices supported, etc.

For each type of independently accessible functionality by a controlling agent, device objects 66a-66N are created in the I/O directory structure representing the different functions that are available that will be accessed by user mode clients. These typically represent filters or other portions of functionality independently available. The driver object 64 and the device objects 66a-66N are created upon installation of the driver code as represented by the enclosing box 68.

Historically, a device object exists for each element of physical hardware. The flexibility in modern I/O systems, however, allows a device object to represent a filter implemented entirely in software. As such, device objects may be readily created for each instance of a filter implemented solely in software. A software filter may therefore be implemented so that each instance as represented by a device object has a one-to-one correspondence with a device object or a single device object may follow the more traditional approach and manage multiple file objects, with each file object representing a client instance of the filter.

Upon a device object, as shown for device object 66a, file objects are created representing independent instances of the functionality represented by device object. While a device object represents a filter and may manage multiple instances of that filter, a file object represents the actual instance of that filter used by a particular entity. Therefore, file object 70 is an instance of the filter defined by device object 66a.

To use a filter, the controlling agent or other user mode client opens a file on a device available in the I/O directory structure. A file object with appropriate context information will be created and a handle to that file object returned to the user mode client. While file objects may be hierarchally related by specifying a "parent" file object during creation, file objects will also have a sibling relationship in that they are all children of the same device object.

Context information within a file object consists of information to manage the I/O interface with user mode clients, the "state" of the entity that the file object represents, etc. The context information has system required information and further includes user definable areas that can be given special meaning. An example of how the user definable area can be used will be shown hereafter discussing the implementation of a validation and IRP routing method.

In order to provide connection pin instances, the file object 70 representing a filter instance will be used as a parent in creating children file objects representing the connection pin instances for a particular filter. While file object 70 will be queried for the connection pin factory definitions and availability, actual file objects will be created for each instance of such a pin factory, using the particular file object as the appropriate informational context in order to validly and correctly create the connection pin instance. For example, file objects 72 and 74 represent connection pin instances for the filter represented by file object 70 and are hierarchally related to file object 70. The connection pin instances, as represented by file object 72 and 74, respectively, may be a data path into and then out of the filter instance (represented by file object 70) which can be used for connecting to other connection pin instances in forming a series of chained filters or other driver functionality.

Just as a pin instance is represented by a file object having a hierarchial relationship to another file object representing the filter instance in order to provide context information for the pin instance, other file objects may be hierarchally related to a pin instance in order to represent other functionality so that proper context information is available. Context information is necessary to distinguish one pin instance from another according to the individual parameters used in creation, such as pin data format, communication type, etc.

Other operational entities, such as a buffer allocation mechanism, a timing mechanism, etc., requiring either an individual context or user mode control through a handle may also be represented by file objects. Furthermore, hierarchical relationships between the file objects (e.g., a buffer allocation mechanism associated with a particular connection pin instance) may be established if necessary by specifying a parent file object during creation of the child file object. These parent/child relationships exist to determine relationship and structure between the file objects representing the operational entities. Additionally, a particular type of "parent" file object will only be able to produce certain types

of "children" file objects, thus requiring the creation validation mechanisms as explained hereafter. Again, such file objects have corresponding handles available in user mode that are returned to a client through a system API call such as `NtCreateFile`.

The handles to file objects are used by user mode clients, such as a controlling agent, to communicate with the kernel mode drivers. The hierarchical chain of file objects, device objects, and driver objects allows the I/O subsystem to traverse back to the driver object through the hierarchically related file objects and device objects to arrive at the entry points into the actual driver code. Such entry points are references (e.g., pointers) to functions in the software driver code. Furthermore, each of the objects in the object pathway between a particular file object and the driver object having the entry points to the software driver code provides important context information for the I/O subsystem in creating IRPs as well references into data structures used for properly routing IRPs according to the routing and validation mechanism explained hereafter.

Handles for file objects and other system objects are process-specific and are the means by which a user mode process will communicate with an underlying object. It is interesting to note that multiple handles may be created to reference a single underlying system object, such as a file object. This means that multiple applications may be feeding information to a pin instance as represented by a file object.

One element of information that is important for interconnecting different drivers is the device object stack depth parameter. This will indicate the IRP stack location of a particular driver object. In this manner, a single IRP may be used and passed between interconnected drivers using the I/O manager, thereby providing a performance enhancement over separately creating and sending IRPs between the various interconnected drivers. Alternatively, each driver could create through appropriate I/O manager calls new IRPs for each successive communication and cause each new IRP to be sent to the next driver in the chain of interconnected drivers.

Referring now to Figures 4A-4C, extensions to the system driver objects, device objects, and file objects are shown that allow validation of file object creation of differing types as well as I/O Request Packet (IRP) routing to appropriate handlers. Figure 4A shows a driver object 76 representing the executable code implementing one or more filters or other driver functionality. Within the driver object, the Windows NT architecture requires a reference to a create handler provided by the software driver developer. According to this embodiment, a multiplexing dispatch function 78 is referenced from the driver object 76 as the create handler and will be used to route messages to particular create handlers depending on the file object type to be created. Operation of the multiplexing dispatch function 78 will be explained in connection with the flow chart shown in Figure 6 hereinafter.

In like manner, other handlers from the driver object will indicate a multiplexing dispatch function and, depending on implementation, they may be the same function. In other words, as explained in more detail below, each type of I/O handler reference (e.g., read, write, device control, etc.) will point to a multiplexing dispatch function that uses the extension data in a device object and the context information in a file object in order to route the message to the appropriate handler. The extension data in the device object that references a validation table will be used when no parent file object is specified on a create operation. Otherwise, the parent file object context information will indicate the correct validation table.

Figure 4B shows a driver object 80 which has a particular device extension area 82 that can be utilized as desired by the driver developer and includes driver specific information. At a defined location within the device extension area 82 of the driver object 80 is a reference to a data structure, known as a file type validation table 84, containing string representations of file object types 86 and references to the associated create handlers 88 for each file type represented. The create multiplexing dispatch function will utilize file type validation table 84 to validate the file object type to be created and then turn control over to the appropriate create handler as will be explained in detail hereafter in connection with the discussion of Figure 6. The string to be validated is found in the IRP create request and originates from the file name string used with the `NtCreateFile` function call in user mode. The `NtCreateFile` call is made within the user mode function cell to create a pin instance or other mechanism.

Figure 4C shows a file object 90 having a file context area 92 that is free to be used by the software driver developer. Reference is made from the file context area 92 to an IRP request handler table 94. The different types of IRP request 96 are associated with references to particular handlers 98, and the appropriate multiplexing dispatch function will use this information to access the correct handler. In the case of determining the correct create handler, a data structure known as a file type validation table 100 is referenced having string representations of file object types 102 and references 104 to the associated create handlers for each file type represented. For children file objects (i.e., file objects that have another file object rather than a device object as parent), the type is represented by a string that is compared to the strings in the file object types 102. When a match is found, the associated create handler is accessed using a reference from the references 104 that is associated with the matched file object type string. If no match is found, then the request is invalid and an error indication made.

Referring now to Figure 5, the installation procedure for setting up the creation validation and mechanism is shown. At step 106, the installing program will make reference in the driver object to the appropriate multiplexing dispatch functions. As shown in Figure 4A, the create handler points to a generic multiplexing dispatch function. In like manner, all

other handler references in the driver object 76 would point to other generic multiplexing dispatch functions germane to the particular handler as necessary. Alternatively, each handler reference could point to the same multiplexing dispatch function that could in turn process the IRP request and route it to the appropriate handler. Such an alternative multiplexing function will necessarily be more complex in order to accommodate different kinds of request (e.g., create, write, etc.).

Next, at step 108, each device object created as part of the software driver executable code installation will be adjusted to reference the file type validation table 84 as shown in Figure 4B. Finally, at step 110, the processing of IRP requests will begin with the multiplexing dispatch function using the file type validation table 84 as referenced from the appropriate device object 80

When a file object is created, the appropriate IRP dispatch table 94 will be created and referenced along with the indexed file object type validation table 100 as necessary. The creation of the file object type validation tables occurs within the provided create handlers according to file object type. The data structures are created representing the IRP dispatch table 94 and the file object type validation table 100 and a reference thereto stored at a specific location with the file context information 92 of the particular file object 90 being created.

Referring now to Figure 6, a flow chart is presented showing the operation of the create multiplexing dispatch function and its validation mechanism including its interaction with the data structures referenced from the system driver objects, device objects, and file objects. At step 112, a user mode process sends an I/O request for creating a file object. This I/O create request is made using an invocation to the system API for NtCreateFile. At step 114, the I/O manager sends the IRP to the multiplexing dispatch function 78 based on the reference in the driver object 76 (see Figure 4A).

Once the multiplexing dispatch function 78 has the IRP for the create request, a test is made at step 116 to determine if there is a parent file object. The information necessary to make this determination will be found within the IRP itself and originally be supplied by the user mode process. The user mode process will supply a handle referencing the "parent" file object as part of the create request and the NT system will create the IRP having the correct reference to the "parent" file object.

If there is no parent file object, the right branch is taken, and the multiplexing dispatch function 78 uses the device extension 82 from the appropriate device object 80 to reference a file type validation table 84 (see Figure 4B) at step 118. Using the validation table 84, the multiplexing dispatch function 78 will validate the file object type at step 120 by comparing the string in the request with the file object types 86 strings.

If there is a matching string as determined at step 122, the appropriate create handler is accessed at step 124. Otherwise the create request is rejected at step 126. The create handler as accessed at step 124 will create, or cause to be created, the file object at step 126. With the created file object, the appropriate create handle will make the file object reference in the file context 92 to an IRP dispatch table 94 that it has previously created.

Again at step 116, it may be determined that there is a parent file object present. If a parent file object is present, as determined at step 116 as found in the IRP associated with the create request, the multiplexing dispatch function 78 uses the file context 92 from the parent file object 90 to reference an IRP dispatch table 94 (see Figure 4C) at step 130. For a create request, the multiplexing dispatch function 78 will access a file type validation table 100, at step 132. Using the file type validation table 100, the multiplexing dispatch function 78 will validate the file object type at step 133 by comparing the string in the request with the file object types 102 strings, as was done above.

If there is a matching string as determined at step 134, the appropriate create handler is accessed at step 138. Otherwise the create request is rejected at step 136. With the appropriate create handler, the file object is created at 140, and the create handler will make a new IRP dispatch table 94 for the newly created file object and will make reference in the newly created file object 90 file context area 92 to the newly created IRP dispatch table 94 at step 142. Note that the same file object structure as shown in Figure 4C is used to explain interaction with both the parent file object and the validly created child file object. While the same structure exists in both cases (once the new file object is created), they will be used differently and contain different information.

Whenever a connection pin instance is created, a connection pin ID is passed that identifies the pin factory in the filter that "supports" the creation of the pin instance. Those skilled in the art will note that the connection pin ID may also be validated as a string in a validation table in much the same manner as the file object is validated and that other implementation variations exist.

In order to make connections between different drivers, a common mechanism must be present to assure that a given driver supports such interconnections. This common mechanism must allow discovery of filter capabilities including connection pin factory capabilities. Furthermore, such a mechanism should also be extensible to provide additional flexibility to driver developers.

One mechanism chosen in the present embodiment for defining compliant drivers and allowing discovery of capabilities are identified "sets" of related items. This is a convenient mechanism to be used with existing I/O communication mechanisms. A set is logically defined as having a GUID (globally unique identifier) to identify the set as a whole and a RUID (relatively unique identifier, e.g., relative within the set itself) for each element of functionality within the set. The

set identifier and any other data structures necessary for operating with the chosen RUID item are passed as part of an I/O control call using the filter handle as a parameter. Only a small number of IOCTLs need to be allocated in order to implement a full system of functionality. As implemented, three different types of sets are established depending on their functions, requiring a total of four IOCTLs. Other implementations may use sets in a different manner. The particular IOCTL will signal the handler for I/O control to interpret or use the chosen element (using the RUID) in a certain manner. Furthermore, control flags may be passed with the GUID and RUID to further specify control information.

The first set type is a property set and is used in connection with values or settings found within the driver or on any associated hardware. Examples of such settings would be transfer rate, volume level, etc. One IOCTL is associated with property sets with a control flag differentiating between a "get" property and a "set" property command. In this manner the same data structure can be used to either set or get a particular property with the driver determining the action required based on the IOCTL used. The correct property is identified by the set identifier consisting of its unique GUID and RUID combination.

Method sets are another type of set used and are a set of actions that can be performed by a driver. Only one IOCTL is needed to identify the method set with the correct method to be actuated identified by the unique GUID and RUID combination for the set identifier. Methods are used to control the driver and include such functions as initializing the driver for use, clearing buffers, etc.

Event sets are used for managing events associated with driver processing, such as device change notification, data starvation notification, etc., or any other notification defined by the set that may be useful to a user mode application. Two IOCTLs are used, one for enabling a specified event and one for disabling a specified event, while any data structures necessary for a given event identified by a RUID can be shared whether enabling or disabling the event.

To use a set, an I/O control operation is initiated using the specified IOCTL and reference to a data structure having the set GUID, RUID, and other necessary data (e.g., control flags, data structures, etc.). For example, setting a volume property on a sound card driver would entail an I/O control operation using a property set IOCTL, a control flag indicating a set property operation, the appropriate GUID for the property set having the volume setting, the specific RUID within that set indicates the volume property, and the new volume setting value.

To query the sets supported, by type, an IOCTL for a particular set type (e.g., property IOCTL, method IOCTL, or event enable IOCTL) having a null GUID and control flags to indicate supported set enumeration are issued as part of an I/O command and a list of set GUIDs supported will be returned. To query supported properties, methods, or events within a given set, the set GUID, set type IOCTL, a null RUID, and control flags indicating enumeration of supported elements are used with the I/O operation. A list of supported RUIDs will be returned as a result of the I/O operation. This will allow a third party agent to determine which, if any, optional elements of an implemented set are supported.

The written specification of a set uniquely identified by a GUID allows a documented mechanism that both driver developers and third party controlling agents may use as an implementation guide. The third party developer will know of a given driver's capabilities based on response to queries and preprogrammed knowledge based on the abstract set definition. Likewise, a driver developer may use the abstract set definition as a guide to implementing a set or group of sets providing known functionality to any third party agent.

In order to provide the connection abilities described herein, a compliant driver must support certain sets. The following tables illustrate some important kinds of information that may be supported in property set format and that can be used in implementing the present invention. The first table refers to properties about a connection pin factory that would be implemented by a filter, while the second table refers to properties about an actual connection pin instance that would be created by using a particular connection pin factory as a template.

TABLE 1

Filter Properties and Their Use	
Property	Description
Connection Pin Factories	Lists the different types of connection pin instances that may be created on a particular filter, each distinguishable type referred to as a pin factory. Note that this is not the total number of connection pin instances which could be instantiated on this device, but the number of unique connection pin types, such as an audio input and audio output.
Connection Instances	Lists the number of instances already created of a given connection pin factory as well as the highest number of instances supported for that particular connection pin factory. If the total cannot be determined until the filter is actually connected, this property will return a -1.

TABLE 1 (continued)

Filter Properties and Their Use	
Property	Description
Data Flow	Lists the possible data flow direction of a connection pin factory with respect to a filter (e.g., into the filter, out of the filter, or either into or out of the filter).
Communication	<p>Lists the communication requirements for a given connection pin factory in terms of processing IRPs. Some connection pin factories may not be interconnected but have other forms of control mechanisms associated therewith such as a "bridge" to a file source for data that represents a source point on a graph. The bridge control mechanism would allow setting of a filename indirectly where information is stored.</p> <p>In an exemplary implementation, an agent (which decides which pin factory to use for making a connection pin instance) must be able to understand the intrinsic meaning of a "none", "sink" or input, "source" or output, "both," and "bridge" communication types for a connection pin factory. For example, a source connection pin instance requires a handle or reference to a sink connection pin instance, etc.</p> <p>In the communication type context, sink and source refer to the disposition of the connection pin instance in processing IRPs. A sink would receive the IRPs for processing, while a source would pass the IRPs onto the next appropriate processing component.</p> <p>There are two communication types that are neither sink nor source and represent end points in the connection graph. An end point represents the place where data either enters or exits from the connected filters. A none designation indicates that the connection type may not be instantiated while a bridge communications type refers to an end point that may be instantiated so that specific properties may be manipulated. For example, a bridge end point that is part of a file reader will likely have a property that will contain the path and file name of a file that stores the data to be processed.</p>
Data Ranges	<p>Lists the possible data ranges that a connection pin factory may support, including the format of the data, if relevant. In one implementation, a count followed by an array of data ranges, which the connection pin type may support, is used as part of the property. In that implementation, if different data ranges are supported under different mediums or interfaces (see below), different connection pin factories are available on a particular filter to accommodate such differences. Furthermore, each data range structure may be extended for format specific detail such as number of bits and channels.</p> <p>The actual data format a connection pin instance uses is set during creation of the instance. The data range property is used to assist in determining what that actual data format should be for a particular connection pin instance and is accessed or queried by a third party controlling agent.</p>
Interfaces	Lists other set GUIDs indicating the supported interfaces on a particular connection pin factory. An interface is the type or types of data that may be communicated through a connection pin factory. For example, MIDI data, CD music, MPEG video, etc., would be interfaces in the sense that data has a particular convention and format that a filter could handle. Such interfaces also comprise protocols for submitting the data. An interface is independent of the medium by which it is communicated.
Mediums	Lists the supported mediums on a particular connection pin factory. A medium is the way or mechanism by which information is transferred, such as IRP-based, sockets, etc. An interface may be defined on top of a variety of different mediums. In the preferred embodiment and implementation explained herein, an IRP-based medium and file IO- based medium is used.

TABLE 1 (continued)

Filter Properties and Their Use	
Property	Description
Data Intersection	<p>Returns the first acceptable or "best" data format produced by a connection pin factory given a list of data ranges. This approach is used to allow a third party agent to determine data requirements when chaining different filters together. In one implementation, the data intersection property is used to determine the best data format produced by a connection pin factory given the constraint of a list of data ranges. The list of data ranges may be acquired using the data ranges property on another pin factory that will be connected as explained previously.</p> <p>A third party controlling agent, which has no knowledge of the data type specifics, may use the data range list of one connection pin factory and return the "best" (e.g., first acceptable data format) data format on the current connection pin factory. Although a set of ranges of the two intersecting connection pin factories could be returned, only the best format is returned by the driver. In this manner, the third party controlling agent can apply this "best" data format to the next driver in the graph in order to create a virtual set of connections before actually initiating the creation of connection pin instances and connecting the entire graph of filters together. This allows the controlling agent to assess the viability of a particular filter graph selected by a user and point out potential problems to the user before actually connecting the graph. The data format returned can also be restricted by the formats available given the connections already made on the filter.</p> <p>This property is capable of returning an error if a particular data format cannot be determined until an actual connection is made or if an intersection is dependent on multiple data formats on different connection points. Essentially, intersection information is provided while the property itself will return a data format.</p>

TABLE 2

Connection Pin Instance Properties and Their Use	
Property	Description
State	<p>Describes the current state of the connection pin instance. Possible states include being stopped, acquiring data, processing data, being paused or idle, etc. The state represents the current mode of the connection pin instance, and determines the current capabilities and resource usage of the driver.</p> <p>The stop state is the initial state of the connection pin instance, and represents the mode of least resource usage. The stop state also represents a point wherein there will be the most latency in data processing in order to arrive at the run state. The acquire state represents the mode at which resources are acquired (such as buffer allocators) though no data may be transferred in this state. The pause state represents the mode of most resource usage and a correspondingly low processing latency to arrive at a run state. Data may be transferred or "prerolled" in this state, though this is not actually a run state. The run state represents a mode where data is actually consumed or produced (i.e., transferred and processed) at a connection pin instance.</p> <p>More resolution in control may be accomplished using custom properties depending on the purpose of the filter and the underlying hardware. For example, in order to make an external laser disk player spin up, one would set some sort of custom "mode" property specific to that class. Setting this property may also change the state of the device but not necessarily, depending on the effect of the mode.</p>

TABLE 2 (continued)

Connection Pin Instance Properties and Their Use	
Property	Description
Priority	<p>Describes the priority of the connection pin instance for receiving access to resources managed by the filter and is used by the filter in resource allocation arbitration. This property, if supported, allows a third party controlling agent to indicate the priority placement of the particular pin instance relative to all other connection pin instances of all other drivers which may share limited resources with this particular connection and instance.</p> <p>This priority property may also be implemented to allow an agent to set finer tuning of the priority within a single class of priority. For example, a priority may have certain subclasses associated therewith. If two drivers competing for the same resources have the same priority class, then the subclass priority is used to determine resource allocation between the two drivers. If the subclass priority is also the same, then arbitrarily, the first connection pin instance will receive the resources.</p>
Data Format	Used to get or set the data format for the connection pin instance.

The previous tables are given by way of example and not by limitation, and those skilled in the art will appreciate that many different properties and schemes may be implemented in order to create the connections between different drivers. One important element is the standardization factor so that different driver manufacturers or development groups may create drivers that may be interconnected since they are able to implement the same property sets.

Another useful property set gives topology information for the internal relationships of input and output connection pin factories on a given filter. This information will state the relationship of input pin factories and corresponding output pin factories on a given filter as well as what type of processing happens between the input and output pin factories. Examples of the processing that occurs would be different data transformations, data decompression, echo cancellation, etc. Such information is useful to an automated filter graph builder that will trace a hypothetical connection path using multiple filters before making actual connection pin instances and connections. Essentially, the topology information explains the internal structure of the filter and exposes this through a property set mechanism to inquiries from third party agents.

Therefore, a compliant driver is simply one that implements the designated property set. This allows a third party controlling agent to make queries and settings to the compliant filter once it is determined that a given property set is supported. The overall goal is to acquire enough information on how to connect the differing filters together in order to form a filter graph.

By using the generic set mechanism, a minimum of functionality may be implemented to support a compliant driver but still allow unlimited extensibility. A set may be defined in a written specification that can be independently coded by a multitude of different driver developers to create a system of interoperable and interconnectable drivers as long as particular sets are implemented. Furthermore, the specification can define mandatory properties, methods, and events that must be supported as well as optional properties, methods, and events that can be implemented depending on the driver functions and advanced capabilities. Besides the basic minimum commonality required, driver developers may incorporate additional functionality by defining their own sets and assigning them a GUID.

Referring now to Figures 7 and 8, an illustration of the process for connecting two kernel mode filters is illustrated. Figure 7 shows a logical block description wherein each filter instance and connection pin instance is represented by file objects. Figure 8 is a flow chart illustrating the steps to creating the file objects and the appropriate connections.

Beginning at step 144, an instance of Filter A 146 and an instance of Filter B 148 are created by a user mode agent. These are created using standard file system API for creating files with a particular device. Filter A 146 and Filter B 148 will be compliant filters or drivers because of their implementing the appropriate property, method, and event sets to support the creation of connection pin instances and for querying the respective filter's capabilities in terms of sets supported and connection pin factories defined for that filter.

The third party controlling agent will then query Filter A 146 and Filter B 148, respectively, at step 150 to determine connection pin factories available and the attributes for connection pin instances that may be created therefrom. These attributes include, as mentioned previously, the connection format and the data format for each individual type of pin instance for each respective filter 146 and 148. The querying will be accomplished using the set based query mechanisms explained in detail previously.

After querying such information, the third party controlling agent will determine the optimal connection format based on the ranges of data formats and connection formats previously queried. This determination occurs at step 152 and places in the third party agent the ability to use the same filters in different ways according to the needs of a selected connection path. The third party controlling agent will use the data intersection property, topology information,

and connection pin factories on both the filters in order to determine how best to select data format and connection arrangements depending on the actual filter graph being made.

Input filter pin instance 154 is created by the third party agent at step 156 using the optimal detection formation determined at step 152. Since input pin instance 154 is a file object, a handle will be returned from the create process that can be used for delivering I/O requests to the input instance 154. Furthermore, the creation of the input pin instance 154 was validated and uses the routing and validity mechanisms shown previously in discussion with Figures 4A-4C, 5, and 6.

In order to finalize the connection, output pin instance 158 is created at step 160 using as a parameter in the NtCreateFile call the handle of the previously created input pin instance 154. The effect of thus creating the output pin instance 158 is to utilize the system file management and I/O management facilities to create an internal IRP stack structure that allows an original write command to be consecutively processed by the variously connected connection pin instances and filters in an appropriate order so as to facilitate direct data flow between the differing filters. This requires that the input pin instance be created prior to the associated output pin instance that will be feeding the input pin instance.

The stack depth parameter of a device object controls how many stack locations are created from IRP sent to this driver. A stack depth parameter is assumed to be one when a device object is initially created and may be modified thereafter depending on the whether multiple drivers are chained together. In the current system, modification occurs, if necessary, when an output pin instance transitions from the initial "stop" state to the "acquire" or other state. Connection pin instance state transition is the mechanism that determines correct stack depth parameter information for proper IRP creation and treatment.

In order to correctly allocate the internal IRP stack structure for a chained set of connection pin instances, it is necessary to transition the connection pin instances out of the stop state in a specified order; beginning with the last input pin instance (in this case input pin instance 154) and working consecutively backwards to an associated (e.g., connected) output pin instance (in this case output pin instance 158). If many filters are chained together, the deepest filter's or bridge's input pin instance must be the beginning point of transitioning and building successively backwards until the initial output pin instance on a bridge or filter is set. In other words, the transition out of the stop state must occur backwards up the chain so that each connection pin instance gets the stack size needed after the previous connection pin instance. Typically, though not necessarily, a connection pin instance transitions from the stop state to the acquire state and for discussion purposes hereinafter transitioning to the acquire state will accomplish the same purpose with respect to stack depth parameter adjustment as transitioning out of the stop state.

Once all pin instances are in the acquire state, stream reads and writes may be issued to the filter graph. It is interesting to note that the system explained herein allows connection of associated input and output pin instances to occur in any order; only the transition from the stop state must occur in bottom up or deepest first fashion. Furthermore, the filter graph is reconfigurable to allow changes to be made after initial creation. When changes are made, state transitions need only occur on those connection pin instances that are in the stop state in order to assure correct stack depth parameter information.

Connection pin factories found on filters represent places where a filter can consume and/or produce data in a particular format. For example, a particular connection pin factory may support a number of different data formats, such as 16 bit 44 kilohertz PCM audio or 8 bit 22 kilohertz PCM audio. As explained previously, the connection pin factories and their different capabilities such as data format can be queried from the filter using the appropriate property set mechanism and the system I/O facilities. Actual connection pin instances are treated based on the information received from the pin factories.

In a streaming environment, where a single stream write or stream read operation from a user mode agent will cause successive processing of the data through the connected filters, two main methods for IRP control can be used as part of the native facilities of the NT operating system. First, a separate IRP may be created by each filter and sent to the next filter for processing which will in turn create a new IRP for further processing down the chain. The other method is to use a single IRP and pass it between the successive filters using standard procedures provided for interacting with the I/O manager. If the first method of creating new IRPs for each successive filter in the chain is used, interconnection order between the filters is unimportant since the filter need only know the destination of the IRP in order to call the I/O manager and send the IRP to the designated filter. If an IRP is reused, it is important that the connection pin instance transitions from the stop state be made beginning from the last filter to receive the reused IRP for processing backwards up to the first filter to receive the reused IRP or to the filter that created the IRP for processing.

The current embodiment and implementation of the interconnected kernel mode filters utilizes IRP sharing advantageously to ease complexity in driver development, allow more robust drivers to be created, and provide more efficient processing. The "bottom up" pin instance state transition path will ensure that the proper stack order is created in the IRP processed by the successive drivers and that each driver object has the appropriate stack depth parameter set. Furthermore, the current state of the receiving input pin instance is checked in order to assure that the state transition sequence has been properly followed. For this reason, the communications property of a particular connection pin fac-

tory will determine the potential flow direction and aid in properly distributing the state transition of connection pin instances.

When creating an output pin instance (or IRP source), a reference to a file object representing an input pin instance (or IRP sink) on another filter will be passed as part of the `NtCreateFile` call. The appropriate create handler will be executed as explained previously using the multiplexing dispatch function and device object/file object hierarchy. This create handler will have access to the device object of the filter having the input pin instance (e.g., Filter B 148 in Figure 7) by way of the input connection pin instance file object (e.g., input pin instance 154). From the device object, the previous stack depth parameter can be read, and the stack depth parameter of the device object for the filter having the output pin instance may be incremented. For example, the device object associated with Filter A 146 will have a stack depth parameter incremented from that of the device object associated with Filter B 148 for the connection illustrated in Figure 7. This normally occurs when transitioning out of the stop state and IRPs are not forwarded while a connection pin instance is in the stop state.

When a filter processes an IRP, it knows which stack frame or location within the IRP stack to access containing information designated for that particular filter by making reference to and using the stack depth parameter of the associated device object. Furthermore, the current filter will prepare the IRP for the next filter in the processing chain by decrementing the device object stack depth parameter to locate the next filters IRP stack location.

The filter code is responsible for preparing the next location in the IRP stack and for calling the I/O manager to pass the IRP to the next filter as designated. In this manner, the filter may designate which file object representing a particular connection pin instance is to receive the IRP and the associated data for processing. Hence, the standard I/O manager calls such as `IoAttachDevice` to stack the respective device objects for sequential processing of IRPs are not used.

It is noteworthy that creating a connection between connection pin instances does not imply creating new device objects to represent the connection. A single underlying device object is used to support an instance of a filter and all connection pin instances on that filter. Specific information necessary for proper data processing is kept within the context area of the file object allowing the context information to be preserved while non-page memory use is kept at a minimum. It is also noteworthy that while an IRP-based medium has been illustrated, other mediums for communication between the interconnected filters may be used, such as direct function calls on non-host hardware-to-hardware communication.

Referring now to Figures 9A-9B and Figure 10, the proper creation, connection, and state transition order of the software drivers as shown in Figure 1 (prior art) and Figure 2 (higher level logical diagram of the interconnected kernel mode drivers) are presented. Figure 9A illustrates the logical structure encompassed by box 162 and the processing steps contained therein. Figure 9B shows the creation of the connection pin instances to complete the interconnection of kernel mode filters and comprises the processing steps encompassed by box 164 on the flow chart shown in Figure 10.

When in the state of Figure 9B, having all interconnections made, the kernel mode filter system is ready for reads and writes in order to effectuate processing. The I/O system will use the IRP stack information properly set by the correct state transition process in order to pass the stream reads and writes onto the differing filter elements by way of their respective connection pin instances. It may be noted that some external software other than the agent used to create the graph, including a bridge or filter itself, as well as hardware will provide data for the stream reads and rights.

After beginning at step 168, the controlling agent 170 will create instances of reader filter 172, decompressor filter 174, effects filter 176, and sound rendering filter 178 at step 180. Furthermore, attachment will be made between reader filter 172 and a disk driver 182 in order to bring the data in from off of the disk drive. Creation of each filter instance is achieved by the user mode controlling agent 170 by using standard I/O calls to open a file on the appropriate device as found in the device I/O directory hierarchy. Such a call will return a handle to a file object representing the instance of each filter.

At step 184, the third party agent will query the effects filter 172, the decompressor filter 174, the effects filter 176, and the sound rendering filter 178 to determine the connection pin factory capabilities. These capabilities include what kinds of input and output pin instances may be created, how many instances of each connection pin factory the particular filter will support, the data format supported on each connection pin factory, the medium or type of communication path, etc. The capabilities are "queried" using the property set mechanism explained in more detail previously and the kernel mode filters are presumed to be compliant to the architecture since they support appropriate "sets" (e.g., property set).

All such query information at step 184 will be used to determine if a chained connection path is possible between the respective filters by creating and connecting the appropriate connection pin instances. The third party agent will determine the types of pin instances needed for interconnection in order to make a filter graph to accomplish a given purpose.

The determination of the connection format based on the supported data formats is determined at step 186. Using the topology information, data format, and data intersection properties on the filter, a hypothetical filter graph may be created. Since connection order is not significant, this need not be done but could save time when trying to build a filter

graph. Should this hypothetical filter graph be created without error, the third party agent will be assured that creating an interconnecting connection pin instances can be done with reliability. Because some queries will return errors unless an actual pin instance is created, it may be necessary to create such connection pin instances before a hypothetical filter graph can be created that will return a reliable indication of viability. Again, the hypothetical filter graph may be tested before any interconnections take place.

Once the correct connection information is known, as determined at step 186, the input pin instances may be created and interconnected and as represented by the loop of processing steps enclosed by box 164 on Figure 10. This loop contains processing steps that will begin at the input pin instance furthest away from the source of the data stream. This last input pin instance is referred to as the "deepest" pin instance and may be created first, followed by the output pin instance associated therewith. A connection, therefore, is the creation of an output pin instance using the handle of a previously created input pin instance.

The pattern continues, with every input pin instance created successively afterwards prior to connection with the associated output pin instance. Such a connection scenario is given by way of example and is not to be limiting of other possible ways of connecting the respective output and input pin instances to form a connection between kernel mode filters according to the present system. The filters may be connected in any order according to implementation as long as the handle from the input pin instance is used during creation of the connected output pin instance on another filter. Furthermore, as explained previously, changes may be made to the filter graph after initial creation (and even use).

In the first iteration of the loop, input pin instance 188 will be created at step 190. After receiving the handle from the create function, the third party controlling agent 170 will use that handle as a parameter in an NtCreateFile call in order to create output pin instance 192 at step 194. By doing this through the first iteration, the sound rendering filter 178 is effectively connected to the effects filter 176 through the corresponding connection pin instances 188 and 192, respectively. In the current implementation, the NtCreateFile call is "wrapped" as part of a function call in an API made available to the user mode clients. This relieves the user mode developer of third party agents from needing to know as much detail and allows all relevant functionality be concentrated in a single user mode API.

At step 196, the third party agent determines if there is any other existing input pin instances to be created. If there are, an input pin instance must be created followed by the corresponding output pin instance on another filter. Eventually, all connections will be made and the third party controlling agent 170 will prepare the filter graph for streamed data processing.

In this fashion, input pin instance 202 will be created on the second iteration of the loop enclosed in box 164 at step 190 while the output pin instance 204 will use the handle of input pin instance 202 as part of its creation at step 194. Finally, on the third and final iteration for this particular example, input pin instance 206 will be created followed by output pin instance 208 to finalize the connection.

At step 197, the third party controlling agent 170 will transition each connection pin instance from the stop state to the acquire state in preparation for streamed data processing through the filter graph. To correctly set the stack depth parameter in each of the device objects for the respective filters, it is necessary to make the state transition beginning with the "deepest" or last connection pin instance (e.g., the last input pin instance to receive data for processing) and sequentially moving "up" the chain of interconnected kernel mode filters until arriving at the first connection pin instance (e.g., the first output pin instance that will provide data into the graph). The first filter or bridge will create the IRP with enough stack locations allocated so that the IRP may be passed successively through each kernel mode filter in the graph in an efficient manner.

Finally, the third part controlling agent 170 issues the stream reads and writes in order to process the data at step 198 before ending at step 200.

As explained previously, each creation of an output pin instance will require the handle of a file object representing the input pin instance to be connected thereto. This file object reference will allow the create handler for the output pin instance to save a reference to the device object corresponding to the input pin instance for current or future access.

More particularly, this allows the stack depth parameter of the device object managing the input pin instance to be accessed by the driver of the output pin instance during state transition from the stop state to the acquire or other state. The value of the stack depth parameter associated with the input pin instance is accessed, incremented, and saved into the stack depth parameter for the device object corresponding to the output pin instance.

The stack depth parameter is used to determine where in the shared IRP stack structure the stack frame information is located for a particular filter and will be different for each filter. By so interconnecting the filters and making the state transition in proper sequence, a single IRP may be passed down the chain of interconnected filters in kernel mode with no necessary communication into user mode.

It may be noted that it is possible to have multiple instances based on the same connection pin factory. For example, an audio mixing filter may mix multiple input pin instances into a single output pin instance in terms of processing. Each input instance is of the same type and the filter may only support one type of input pin. Such an arrangement would also be an example of having multiple inputs to a single output.

The converse is also true wherein a splitter filter may have a single input connection pin instance while providing

multiple output pin instances thereby multiplying streams of data. Those skilled in the art will note that many variations and useful combinations can be made from the connection mechanism explained herein according to actual implementation and the needs thereof.

The uniformity and standardization achieved by requiring all compliant filters to support a common mechanism (e.g., property sets, methods sets, and event sets) that can be independently implemented by driver developers allows a controlling agent to conveniently connect compliant filters provided by various different software providers. Furthermore, many of the facilities in terms of connection pin factories needed in one circumstance may not be needed in another circumstance. A determination of the necessary connection pin instances is made initially by the third party controlling agent that makes the actual interconnections between different filters.

Referring now to Figure 11A, operation of a buffer allocator mechanism is shown as used between multiple processing components. Also shown is the data flow (i.e. actual data transfers between buffer frames) between buffers associated with particular processing components. While control will pass to each processing component, data will be transferred only when necessary with some processing components performing their data manipulations and returning the data to the existing buffer frame. In other words, data may be processed in the same location without being transferred to a new buffer and is said to be processed "in place".

A sink processing component 210 will have a buffer allocator mechanism 212 (represented by a square) as part of its functionality. A buffer allocator mechanism will only be present in processing components where it is necessary to ensure certain placement of the data into specific memory such as on board memory on a sound or video processing card or where the previous buffer has unacceptable characteristics, such as byte alignment, frame size, etc. Furthermore, references to the buffer allocator mechanism to be used for allocating frames of buffer memory is indicated by a circle and all processing components will have such references. Note that the source processing component 214 has a buffer allocator reference 216 that references the sink buffer allocator 212 as represented by arrow 218. Furthermore, the transfer processing component 220 has a null buffer allocator reference 222 and the sink processing component 210 also has a null buffer allocator reference as indicated by the empty circle 223.

In this simple processing example, the source processing component 214 allocates a buffer frame 224a by using the sink buffer allocator 212 and accessed using buffer allocator reference 216. The allocated frame 224a is filled with data by the source processing component 214 as represented by arrow 226. It may be noted that the source processing component may perform some data manipulation or transformation before writing the data into the newly allocated frame 224a.

At this point, the source processing component 214 has finished processing and turns control of processing to the transform processing component 220 as represented by arrow 228. The transform processing component 220 will do no buffer allocation or transferring of data from one buffer to another since no buffer allocator mechanism is indicated since buffer allocator reference 222 reference has a null value. Therefore, the transform processing component 220 performs an "in place" data transformation in the allocated buffer frame 224b as represented by arrow 230.

Since the data has not been transferred to a new buffer frame, buffer frame 224a, frame 224b, and frame 224c are the same frame and are simply passed in succession to different processing components. Arrow 231 represents passing the allocated frame between source processing component 214 and the transform processing component 220.

Finally, the transform processing component transfers control of processing to the sink processing component 210 as represented by arrow 232. It may be noted that along with processing control, the same frame is passed for processing as shown by arrow 234 between the frame 224b and 224c. Again, as depicted herein, frame 224a, frame 224b, and frame 224c are all the same frame allocated originally by source processing component 214 and are shown separately for illustration purposes.

The sink processing component 210 will finish processing the data and free the allocated frame 224c within a buffer as presented by arrow 236. Since the sink component 210 is no longer using the buffer, the arrow 236 points inward to the sink processing component 210 and the frame may now be deallocated or reused.

Figure 11B shows how a buffer allocator mechanism would be logically implemented in the scheme of interconnected kernel mode buffers that has been explained previously. Figures 11A and 11B both represent the same filter graph topology and are used to better illustrate operation of the buffer allocator mechanism. The relevant drivers and portions thereof each have access points that allow user mode clients to control the driver and are represented by file objects. Intercommunication occurs using IRPs (whether created by kernel mode drivers or by the NT executive in response to user mode I/O operations).

An instance of source processing component 238 (represented by a file object) has associated therewith an output pin instance 240 (also represented by a file object) that provides a source of connection to another filter instance. An input pin instance 242 that is a "child" of a transform filter 244 was created having a reference to the output pin instance 240 in the manner explained previously in more detail. In like manner, a sink filter 246 having an input pin instance 248 is connected to an output pin instance 250, that is related to the transform processing component 244.

In the system of interconnected kernel mode software drivers, a buffer allocation mechanism is related to input pin instances and is said to be created or formed on the input pin instance. Furthermore, output pin instances will logically

have reference to a buffer allocation mechanism, if necessary, and the filter having the output pin instance will utilize that reference to make any buffer frame allocations and data transfer to new frames prior to turning control to another processing component. As explained, a null reference will indicate that no data transfer to a new frame is necessary and that processing may occur in the existing frame, (*i.e.*, after processing, the data is returned to the same buffer frame). Whether a buffer allocator reference is present is determined by the initial negotiation of the third party controlling agent that created the filter graph.

The buffer allocator mechanism formed on an input pin instance 248 is represented by a file object while the dashed line 254 shows that the output pin instance 240 has reference (*e.g.*, a pointer or handle) to the file object representing the buffer allocator 252. In the example shown in Figure 11B, frames of memory are allocated from system memory 256.

Since the filters may be interconnected in a number of different ways, a buffer allocator though available, may not be necessary. A file object representing an instance of a buffer allocator will only be created should the third-party controlling agent interconnecting the filters determine that it is necessary. In this manner, filters may flexibly be connected in a variety of configurations and still maintain optimal data transfer characteristics. Furthermore, default system buffer allocators can be provided to further reduce the development effort of driver developers.

A third-party controlling agent will also query buffer requirements of connection pins as part of creating a hypothetical model before making actual filter connections. While some implementations may allow queries before pin instantiation, the present embodiment requires that the actual connection pin instance be created before the buffering requirements can be ascertained. Furthermore, the exemplary embodiment disclosed throughout is queried through use of the set mechanisms explained previously.

When a third party client or controlling agent completes the interconnections of kernel mode filters to make a filter graph, it then initiates negotiation of allocator requirements using the handles of the input pin instances (or sink pin instances). By convention, the input pin instances define buffer allocation requirements and provide a buffer allocation mechanism while the output pin instances have reference to any appropriate buffer allocation mechanisms associated with the input pin instances. Those skilled in the art will understand that other conventions may be used to effectively accomplish the same results including reversing the buffer allocation responsibilities between the input and output pin instances.

Buffer allocation requirements are ascertained by using a particular property set mechanism that will be supported by all compliant filters. It may be noted that the "buffer allocator" property set may be organized in a number of different variations as may any other set. For example, the property set may have a single property wherein the property is a data structure having segmented information or the property set may have multiple properties, one for each distinct framing requirement element. A single property composed of a data structure is more efficient in some circumstances since fewer property set queries are necessary by the third party controlling agent in order to retrieve all the framing requirement information. Furthermore, a single data structure could be used not only to query requirement information but also for specifying parameters when an actual buffer allocator is created.

The table below shows a non-exhaustive list of framing requirement elements that may be included in a data structure or as individual properties. Also included is an explanation of how such an element would be used in an exemplary embodiment.

TABLE 3

Element	Description
Control Options (Create)	<p>This element contains control options that are specified during buffer allocator creation on a given connection pin instance. Some options include:</p> <p><i>System Memory:</i> This control option indicates the use of system memory for buffer frame allocations. When specified, the buffer allocator must allocate memory from the pool (as specified in the allocation pool type element) located in system memory. If this control option is not specified, it is assumed that the input connection pin instance (or sink pin) provides a system address mapping to on-board RAM or other form of storage on a device attached to the system. Such a device may be an add in card, a PCMCIA card, etc.</p> <p><i>Downstream Compatibility:</i> This control option specifies that the allocator framing elements of the buffer allocator being created are compatible with the downstream allocator. This option is normally specified when an in place modifier is assigned an allocator for copy buffers. If the filter is not required to modify a given frame, it may submit the frame to the downstream filter without allocating an additional frame from the downstream allocator.</p>
Requirements (Query)	<p>This element contains the allocator requirements for a connection pin instance (input or sink) that are returned during the query of a connection point. Some potential requirements include:</p> <p><i>In-place:</i> This indicates that a connection pin instance has indicated that the filter may perform an in-place modification. Otherwise, a buffer must be allocated for receiving the modified data.</p> <p><i>System Memory:</i> The connection pin instance (input or sink) requires system memory for all buffer frame allocations. If this requirement is not set, it is then assumed that the connection pin instance provides a system address mapping to on-board RAM or other form of storage on a physical device.</p> <p><i>Frame Integrity:</i> This requirement indicates that the connection pin instance requires that downstream filters maintain the data integrity of specified frames. This may be used when statistics are required of a previous frame in order to properly process the current frame.</p> <p><i>Allocation Required:</i> This requirement mandates that connection point must allocate any frames sent.</p> <p><i>Preference Only Flag:</i> This requirements flag indicates that the other requirements signalled are preferences only. This means that the requirements indicate a preferred way of operation but the connection point will still operate correctly and process frames which do not meet the specified requirements.</p>
Allocation Pool Type	This element determines the allocation pool type from which kernel mode entities will receive system memory for buffer frame allocation.
Outstanding Frames	This element indicates the total number of allowable outstanding buffer frames that can be allocated. In this manner, the total size of the buffer can be controlled. Furthermore, specifying a zero for this element indicates that the filter (through a particular connection pin instance) has no requirement for limiting the number of outstanding frames.
Frame Size	This element specifies the total size of a buffer frame. Again, specifying a zero for this element indicates that the filter (through a particular connection pin instance) has no requirements for this member.

TABLE 3 (continued)

Buffer Allocator Framing Requirement Elements	
Element	Description
File Alignment	This element specifies the alignment specification for allocating frames on certain boundaries. For example, an octal byte alignment may be specified. Furthermore, in the exemplary embodiment, the minimum alignment capability is on quad boundaries providing optimal data access speed on current PC hardware architectures.

Those skilled in the art will appreciate that there may be other relevant framing properties that may be included. Furthermore, the buffer allocation mechanism as described herein functions in substantially the same manner regardless of whether more buffer frame elements than specified in Table 3 are included or whether a subset of those specified are implemented.

When the filter graph requirements have been determined by the third party controlling agent, then the appropriate kernel mode buffer allocators may be created on the appropriate input connection pin instances. The client creates a buffer allocator instance using the handle of the appropriate connection pin instance and by specifying the appropriate creation parameters for the buffer allocator. In this manner, the resulting file object that represents a buffer allocator will be a child of the connection pin instance and will use the context information from that file object and the file object representing the instance of the filter itself for its proper creation.

In other words, the same mechanism explained previously for validating creation of a connection pin instance and for routing messages to the appropriate handlers for a given connection pin instance will apply in like manner to the creation of an instance of a buffer allocator. Again, the NtCreateFile call will be wrapped in a higher level function call as part of an API that third party controlling agent may use.

Buffer allocator instances will be created, if at all, only on input pin instances in the presently explained embodiment.

If a buffer allocator instance handle is not specified to an output connection pin instance handle through an appropriate API call, then the filter can assume that the stream segment submitted via stream I/O controls meet the filter's requirements and it may freely modify the data in-place.

A system supplied default allocator may be used by filter developers to simplify the task of providing buffer allocation capabilities to input connection pin instances. The default allocator provides for a system memory buffer frame allocation for those device drivers that are able to transfer data from system memory but require specific memory allocation properties. By using the default buffer allocator, a filter developer is relieved from the task of actually providing code to do the buffer allocation. The filter will still be written, however, to support the buffer allocation requirements request by supporting an appropriate property set.

To use the default allocator, the filter designer will (1) provide code to respond to the buffer allocation requirements requests, and (2) place a default allocator creation handler reference in the validation table of the particular connection pin instance to which the default allocator pertains. In other words, when a create request for the allocator type mechanism is submitted through the filter, a specific GUID string is matched in the validation table that in turn allows access to the default allocator creation handler.

The default buffer allocator uses system memory and will operate in accordance with the buffer allocator framing properties submitted as part of the creation request. Such a default allocator is likely to be used by various transform filters depending upon their particular function and interconnection order.

Filters requiring buffer allocators for on-board memory or other device dependent storage methods can provide a specific buffer allocator by supporting a buffer allocator property set and method set. For a filter specific allocator, the filter will be responsible for providing program code to implement the entire functionality. The operation of the buffer allocator, however, will be the same as for every buffer allocator, whether default or filter specific, so third party agents may interconnect the filters and buffer allocators properly.

The file object representing the buffer allocator, when successfully created, has contained in the file context area a pointer to a data structure. This data structure will contain a dispatch table for directing IRPs to designated handlers based on the various IRP codes (e.g. create, I/O control, etc.) as well as other data areas and structures for maintaining the state of the buffer allocator. Some implementation dependent information that may be tracked includes reference to the file object of the connection pin instance to which the buffer allocator pertains, reference to an allocation framing requirements data structure, an event queue of those clients waiting for events, a queue of those allocation requests (received by an IRP, for example) that are outstanding, etc.

There are two interfaces, or ways of communicating, available to the buffer allocation mechanism disclosed herein in the exemplary embodiment. First, all allocators must provide the IRP-based interface in order to communicate properly with user mode clients. Optionally, if the allocation pool type can be serviced at the dispatch level (a raised priority

level at which a limited subset of services are available, but at which lower priority tasks are blocked out on that processor) of the operating system, a function table interface may be supported allowing interconnected filters to use direct function calls (during DPC processing) for higher performance. This allows event notifications to be communicated between the interconnected filters without the extra overhead of passing an IRP through the I/O manager, or scheduling a thread of execution and waiting for a context switch.

The IRP-based interface will successively process IRPs in the following fashion. When an allocate request is submitted, the buffer allocator will complete the request and return a pointer to an allocated buffer frame or if all frames have been allocated, the allocator marks the IRP pending and adds the IRP to the allocator's request queue for processing, in approximately FIFO order. Finally, the buffer allocator will return status pending to the caller.

When a buffer frame is made available to the allocator, the allocator attempts to complete the first request in the request queue for the IRP-based interface. Those IRPs that could not be serviced previously will be in this request queue and will be waiting for the newly freed frame for processing. Otherwise, if no work awaits in the request queue, the frame is returned to the free list.

The dispatch level interface using the direct function call table operates in the following fashion. When an allocate request is submitted by making a function call (which may be done during a DPC), the allocator returns a pointer to the frame if one is available or returns null immediately. The kernel mode requester may then wait for a free event notification to know that there is a free frame available. Upon receipt of this notification, the kernel mode requester will attempt the allocate request again.

Note that it is possible for both the dispatch level interface and the IRP-based interface to contend for an available free frame. Also, if there are allocation request IRPs waiting to be completed in the request queue, the allocator must schedule a worker item if the current IRQL is not at the passive level when the caller frees a frame since using the direct call interface implies the possibility of being at the DPC level. Essentially, the waiting queue of IRPs does not find out about the free frame until the worker item has been run.

Furthermore, the buffer allocation mechanism explained herein is adaptable to be used with MDLs (memory descriptor lists) in order to further reduce inter buffer transfers. Such an implementation would allow greater seamless interaction with the system memory facilities of the NT operating system.

Referring now to Figure 12, the interconnected filter system of Figures 9A and 9B are shown utilizing the buffer allocation mechanism disclosed previously. Once the controlling agent 170 has made interconnections between the filters, it will query each input pin instance for its buffer requirements. As shown herein, the sound rendering filter 178 will need to allocate buffer frames from the sound card memory 258 and the decompressor filter 174 will allocate memory from system memory 260 which may receive the data directly from disk drive 262.

Controlling agent 170 will create a buffer allocator 264 that is represented by a file object and being formed on the input pin instance 188. The buffer allocator 264 will allocate buffer frames from the sound card memory 258 and a reference to the buffer allocator 264 will be set in the output pin instance 204 as represented by dashed line 266. This reference will be the handle to the file object representing buffer allocator 264 and will be used by the decompressor filter 174 to allocate buffer frames as necessary before transferring data to the new buffers.

In like manner, buffer allocator 268 is also represented by a file object and created or formed on input pin instance 206. This buffer allocator 268 will manage allocation of system memory 260. The controlling agent 170, after creating buffer allocator 268, will store a reference thereof in output pin instance 208 as represented by dashed line 270. Again, buffer allocator 268 will be responsible for allocating buffer frames between system memory 260 so that data may be transferred therein from disk 262.

The controlling agent will also put a value of null for the buffer allocator reference of output pin instance 192 to thereby indicate that an in place transformation will occur wherein the effects filter 176 will read the data from the existing buffer in sound card memory 258 and replace the data back into the sound card memory 258 after applying whatever transform or effect is required. Alternatively, if the controlling agent does not set the buffer allocator reference, it will be assumed to have a value of null and an in place transformation will occur.

Referring now to the flow chart of Figure 13 and the logical diagram of Figure 12, operation of the buffer allocators will be shown. This process occurs after the interconnections are made and the stream reads and writes are given from the controlling agent 170 to the reader filter 172.

Initially, the reader filter 172 allocates a frame in the system memory using the decompressor filter 174 buffer allocator 268 at step 272. The output pin instance 208 of reader filter 172 will have a handle to the file object representing the buffer allocator 268 as represented by line 270, and will therefore have direct access to controlling the buffer allocator 268.

Once the file reader filter 172 has access to an actual buffer frame in system memory 260, it will fill the frame with data from disk 262 as represented by arrow 274 at step 276. It may be noted that the reader filter 172 may make a transformation or other manipulation on the data as it brings it into system memory 260.

The file reader filter 172 then initiates a stream write to the decompressor filter 174 at step 278. This stream write will be passed by way of an IRP through the NT I/O manager. At step 280, the decompressor filter 174 will allocate a

frame of sound card memory 258 using the buffer allocator 264. The decompressor filter 174 will have knowledge of the buffer allocator 264 since a handle thereto was stored with respect to output pin instance 204.

The decompressor filter 174 will decompress the data and transfer the data to the frame previously allocated in sound card memory 258 as represented by arrow 284. Note that in decompressing the data there may be more frames allocated from sound card memory than exist in system memory. The extra ratio of buffer frames is necessary to accommodate the decompression effect on the data.

It is important to note that when transferring data from one buffer to another buffer that there may not be a 1:1 correspondence in terms of the amount of data transferred. In other words, the receiving buffer may require more or less space (or frames) depending on the nature of the filtering or transformation that takes place between the buffer transfer.

Once the decompressor filter 174 is finished decompressing a particular frame, it passes the stream write to the effects filter 176 using the facilities of the NT I/O manager. The effects filter 176 receives the stream write IRP at step 288 and processes the data in place in the existing sound card memory 258. Such effects processing would generally correspond to a 1:1 replacement of data requiring neither more or less buffer memory.

Once the effects filter 176 is finished processing the data in place, the stream write IRP is passed to the sound rendering filter at step 290. Again, the mechanism causing the stream write IRP to be transferred to the sound rendering filter 178 is the NT I/O manager facilities called by the effects filter 176.

Finally, the sound rendering filter 178 will receive the stream write IRP at step 292 and control the sound card hardware in order to form the actual rendering of the sound data existing in sound card memory 258. At that point, the sound card buffer frames that had been previously allocated can be reused to service writing requests or freed if there are no outstanding requests. The availability of a buffer frame is made known to the decompressor filter 174 so that any waiting stream writes may be used to process and place data into the freely allocated buffer frames. In like manner, the buffer frames of system memory 260 are either reused or freed.

Referring now to Figure 14, an overview system is shown wherein a master clock mechanism can be used to synchronize two data streams that are generated in real time and said to be "live." An audio renderer 296 receives a live audio stream 298 for processing as represented by arrow 300. As the audio renderer 296 processes the data, it is perceived as sound through a speaker 302 once the audio renderer 296 has converted the data samples into the appropriate electronic signals and sends the signals to the speaker 302 as represented by arrow 304.

In like manner, a video renderer 306 receives a live video stream 308 as represented by arrow 310 for processing. As a result of being processed by the video renderer 306, the data samples are converted into appropriate video signals and are transmitted to a monitor 312 as represented by arrow 314 where they may be perceived as images.

Typically, the live video stream 308 and the live audio stream 298 will come simultaneously into the rendering components. Since it is important to match the video with the audio when rendered at the speaker 302 and the monitor 312 in order to produce a coherent presentation, processing of the live video stream 308 and the live audio stream 298 must be synchronized.

In order to accomplish this, one data stream is chosen to be the master. For purposes of this disclosure and explanation, the live audio stream 298 has been arbitrarily chosen to be the master data stream. Note that nothing prevents the live video stream 308 from being chosen as the master or any other stream of data for that matter. Furthermore, an external clock, based on a data stream entirely external to the filter graph, may be used to synchronize both the live audio stream 298 and the live video stream 308. Therefore, the master clock mechanism 316 is shown appended with a solid line to the audio renderer 296 and has a dashed line 318 to the video renderer indicating that the video renderer will in some manner receive synchronization information from the master clock 316 that is using the audio stream as the master.

The master clock 316 may also be used by the audio renderer 316 in order to match the rate of processing at the audio renderer 296 with the rate of the original audio stream 298. In other words, the live audio stream 298 was produced by another processing component, such as digitizing filter/digitizing hardware on another system. Since the audio stream is composed of digitized samples, the samples were produced at a rate according to the hardware and oscillator of an original processing component. This origination rate for the original processing component to produce the live audio stream 298 is found in or derived from the time interval information of the live audio stream 298.

Sample information in a data stream may contain time interval information in a variety of ways. One way is for each sample or group of samples to have a timestamp associated therewith. Furthermore, the format of the data itself may imply the sample rate. For example, by the definition of the data format, each sample may be taken at a certain time interval (e.g., 1/30th of a second for a video sample or 1/22,050th of a second for audio data). Additionally, a timestamp may indicate a presentation time that is matched to a clock at the rendering processing unit to determine when a particular sample is to be rendered or "presented." Also, a timestamp may indicate some other time value with the relative difference between time stamps giving time interval information. Finally, mixtures of data format and timestamp information may be used in order to provide the time interval information necessary for creating a time clock based on the position within a reference data stream being processed.

It may be noted that "discontinuities" may be coded into the data stream that may be based on time or data gaps.

A time discontinuity is where the stream simply jumps to another presentation time, such as when recording ended at a certain time and suddenly began at a later time. A data discontinuity is where the stream may have periods of repeated data that can be more efficiently coded in the stream, such as representing periods of silence by a given code followed by duration of the silence.

One time value provided by the master clock 316 is based on the time interval information found in a stream of data samples and known as media time. Media time is positional since it represents a position within the master stream, which in this case has been chosen as the live audio stream 298. Furthermore, when processing quits at the audio renderer 296, the media time is also frozen at the master clock 316. A discontinuity in the data stream will be processed to that "position" within the stream continues to advance while in the discontinuity by using the underlying hardware oscillator or other source to track the advancement of time. In this manner, multiple data streams may still be synchronized to a data stream having discontinuities of time or data therein.

By synchronizing the video renderer 306 to the media time found on the master clock 316, which in turn is based on the time interval information in the live audio stream 298, no overhead is necessary to manage stopped or paused states since media time advances ("ticks") only when the live audio stream 298 is being processed at the audio renderer 296.

It should be noted that the video stream 308 may also need to be rate matched as well as synchronized. The rate matching of the video stream 308 would occur in much the same manner as will be explained hereafter in connection with the live audio stream 298 and the audio renderer 296. Naturally, the live video stream 308 rate matching is based on the hardware oscillator (physical time) of the video renderer 296.

Also available at or otherwise provided by the master clock 316 is a physical time value that is based on the underlying hardware in the audio renderer 296. Physical time usually progresses regardless of whether the audio renderer 296 is actively processing audio data or is in a stopped or paused state. Furthermore, by sampling the rate of progression of the media time (representing the originating hardware processing rate) and the physical time (representing the actual processing rate at the audio renderer 296) adjustments may be made at the audio renderer to process the live audio data 298 so that the rates are relatively matched.

While rate matching is not an important issue when the rate of data production is under control of the processor such as reading stream data from a local disk or other storage device, it is crucial when processing a "live" data stream that continuously produces data and the production rate is out of the control of the processor. As used herein, a "live" data or media stream refers to a stream produced at a rate that is not controlled by the processing components and, therefore, must be processed completely in real time or lost. Examples of a live data stream include but are not limited to the following: a real time feed where the stream is produced in real time either locally or over a network connection, receiving the stream from a network connection regardless of whether it is a real time feed or previously stored data sent over the network, etc. In all cases, the stream must be processed as it is received.

If the live media stream produces too much data for the audio renderer 296 to process, the audio renderer 296 will eventually lose data due to finite buffering capabilities in a condition that is known as data "flooding." On the other hand, if the live audio stream 298 produces data at a lower rate than the audio renderer 296, the audio renderer will suffer data "starvation" in that it will not have data to process. In either case, it is beneficial for quality rendering of the live audio stream 298 to make adjustments to match the rendering rate at the audio renderer 296 with the origination rate as represented in the time interval information in the live audio stream 298.

One way of making rate correcting adjustments is to periodically duplicate data samples so that the audio renderer will not become starved or to throw away samples so that the audio renderer will catch up with the live audio stream 298 that outproduces its capabilities. In either case, duplication or omission of data samples would judiciously occur at intervals over stream processing so that a minimal amount of impact is perceived at the speaker 302.

Rather than synchronize directly to the master clock 316, it may be useful to translate another clock's time base to match the media time found in master clock 316 or to translate presentation time information in another media stream to match the media time of the master clock 316. In order to facilitate such translation, the master clock 316 will provide a correlated media time value that consists of two separate time values, the media time value and a reference time value common to other components, and does so in an atomic (*e.g.*, unitary) operation to reduce the introduction of timing errors. A correlated physical time value is also provided.

The reference time is typically based on the PC clock which should be available to all processing components or other clock mechanisms found on the system though another clock could be chosen. The atomic operation may make use of native system facilities for implementation so that both values may be retrieved with the least amount of interruption or other interference. The correlated time (either media or physical) can be used by another processing component, such as a filter, to determine a relative offset or delta between the media time and the reference time in order to make any necessary translation with a minimum of error introduced as will be explained in more detail hereafter.

Referring now to Figures 15A and 15B, a clock mechanism as described above is implemented with the interconnected filter system explained previously. Figure 15A represents a live video stream processing set of filters that is synchronized to an audio stream and receives synchronization information from a master clock mechanism. Figure 15B is

a set of interconnected filters for rendering a live audio stream that is used to generate the master time reference for a master clock mechanism.

Referring first to the video stream processing components in Figure 15A, a live video stream 320 is received by a video reader filter 322 as represented by arrow 324. The source of the live video stream 320 may be from on-board digitizing hardware at the computer, packets of digitized samples received over a computer network, etc. The live video stream 320 is comprised of digitized samples that may be rendered into actual video images and the samples have time interval information associated therewith as explained previously. Again, according to the interconnected filter system explained herein, the video reader filter 322 is represented by a file object.

The video reader filter 322 has an output pin instance 326 that is connected to an input pin instance 328 on a video decompressor filter 330. Video decompressor filter 330 will receive the live video data in compressed format and decompress it into a format acceptable to video rendering hardware. After processing frames of video data, the decompressor filter will turn control over to the video rendering filter 332. A connection is made between the video renderer filter 332 and the video decompressor filter 330 by means of the output pin instance 334 and the input pin instance 336, respectively.

The input pin instance 336 will receive timing indications for synchronization from the master clock mechanism in the audio stream processing components (see Figure 15B) as represented by box 338 and dashed line 340. This synchronization timing information may be received in a variety of different ways. First, a controlling agent 342 that interconnects the filter graph topology shown in Figures 15A and 15B may have a clock mechanism 344 make event notifications based on either a specific position (*i.e.*, time) within the master audio stream or a specific time interval may be chosen for interval event notifications. Alternatively, the video rendering filter 332 may query the clock mechanism 344 (see Figure 15B). Because the clock mechanism 344 is represented by a file object, the controlling agent 342 can make appropriate references to the clock mechanism 344 to the input pin instance 336 or the video renderer filter 332, thereby allowing the video renderer filter 332 to query the clock mechanism for correlated time values in order to make translations, if necessary, or media time values to assure that processing is synchronized.

Finally, the video renderer filter 332 controls the video hardware 346 and transfers data thereto as represented by arrow 348. The video hardware will render the data and generate video signals causing images to be displayed on monitor 350.

Referring now to Figure 15B, the processing components of the live audio stream are now explained in detail. Again, the same controlling agent 342 will create the respective filter instances along with any necessary input pin instances and output pin instances in order to make the interconnections between the respective filters and it is the controlling agent that makes the entire kernel mode filter graph topology.

The live audio stream 352 is initially brought into the filter graph topology by the audio reader filter 354 as represented by arrow 356. The data is decompressed by the audio decompressor filter 358 that is connected to the audio reader filter using respective output pin instance 360 and input pin instance 362. Again, each filter and connection pin instance described herein is represented by a file object and created as explained previously.

The audio decompressor filter is connected to the audio renderer filter 364 through the output pin instance 366 and the input pin instance 368, respectively. Also associated with the input pin instance 368 is an instance of a clock mechanism 344 that will provide or send indications of media time, physical time, and/or correlated times to other entities, particularly the input pin instance 336 for the video renderer filter 332 (see Figure 15A) as represented by box 370 and dashed line 372. The actual implementation and availability of the clock mechanism 344 will be explained in more detail hereafter. By using the clock mechanism 344, other data streams may synchronize processing to the processing of the "master" or "reference" audio stream 352 as illustrated in Figure 15B.

The audio renderer filter 364 controls the audio hardware 374 and will pass control information and data as represented by arrow 376. The audio hardware, under the control of the audio renderer filter 364, will render the digitized data samples into real time electronic signals that may be transmitted to and perceived from a speaker 378.

When creating interconnected kernel mode filters as explained previously, each filter instance is created from a "device" available to the system. In this manner, the video reader filter 322 and the audio reader filter 354 may be separate instances of the same file reader device.

Additionally, because a particular filter may support many connection pin factories for creating connection pin instances, the video decompressor filter 330 and the audio decompressor may be instances of the same decompressor device. When the controlling agent creates connection pin instances on the respective filter instances, different pin factories are chosen as templates for creating the connection pin instances depending on the type of data being processed through the filter. This is the inherent flexibility in the interconnectable kernel mode filter system disclosed herein.

In order to form clock mechanism 344, a controlling agent 342 will typically query a connection pin instance, such as input pin instance 368 on the audio renderer filter 364, to determine if a clock mechanism is available. Other implementations may require the controlling agent 342 to simply attempt creation of a clock mechanism on a connection pin instance, thereby determining clock mechanism availability by trial and error. Furthermore, a clock mechanism may exist in user mode and send notifications, etc., to input pin instances of filters in kernel mode through a kernel mode

proxy of that clock or some other mechanism.

By chosen convention in the exemplary embodiment, timing event notifications and clock mechanisms are typically found in association with input connection pin instances. Clearly, however, any entity represented as a file object may receive event requests by way of IRPs or by a direct procedure call after retrieving the direct procedure call table through a previous IRP request.

A system supplied default implementation of a clock mechanism may be supplied or made available to software filter developers in order to provide timing functionality without requiring additional code development. Formation of any clock mechanism (supplied default or custom) occurs by specifying a GUID string value in the create request that will be used in a connection pin instance validation table in order to access the correct creation handler for forming the clock mechanism on the connection pin instance. The process of file object creation, validation, and designation of IRPs was discussed in more detail previously in connection with Figures 3, 4A-4C, 5, and 6. In the same fashion as creating a connection pin instance, a clock mechanism create request may be properly routed and validated.

In order to use the default clock supplied by the system, a driver developer will call the default clock create method from within the create handler that they create and provide to the system. The create handler is accessed from the context area of a file object representing a connection pin instance in the manner described previously. The driver developer, therefore, implements a driver specific clock mechanism, and the appropriate driver specific create handler would be placed in the validation table. Complexity of create handler development is reduced by making a call to the default clock create method rather than coding and implementing the clock functionality.

For a clock mechanism to be made according to the interconnected filter system shown throughout, it must support a certain property set and event set. Naturally, the default clock mechanism supports the specified sets and any filter specific clock mechanism would require the filter developer to write and provide code to implement the clock mechanism property set and event set. In this manner, any controlling agent having a handle to a clock mechanism may query for support of the specified set and know how to operate the clock mechanism. Table 4, below, indicates a number of properties that may be included in a clock mechanism property set.

TABLE 4

Clock Mechanism Properties	
Properties	Description
Media Time	This property returns the current media time on the clock as represented in standard units. For example, in the exemplary embodiment, this would be timed in 100 nanosecond units. In the common case where the clock mechanism is being presented by some processing component, such as a filter or connection pin instance on a filter, which is processing a timestamped data or control stream, the media time value returned by this property reflects the current position in the stream. Furthermore, any data or control stream having time interval information may be used to allow the clock to properly operate. If the underlying processing component becomes data starved, the media time as presented by this property also stalls. Furthermore, if the data stream or control stream is comprised of timestamped data, and the timestamps on the stream are modified, the media time as presented by this property will also reflect this modification. Therefore, if timestamp information is "translated" from one time frame to another, any clock mechanisms based on that data stream will also be translated. Finally, if the underlying component is directed to stop or pause processing the data or control stream, the media time as reflected by this property will also stop.

TABLE 4 (continued)

Clock Mechanism Properties	
Properties	Description
Physical Time	This property returns the current physical time for the clock mechanism. The physical time is a continuously running time which is based on some underlying hardware oscillator. For the default clock mechanism, this would be the PC clock. For a filter specific clock mechanism, it may be the actual rendering hardware clock. Physical time does not stall when stream starvation occurs and may not stop when stream processing is stopped or paused by the underlying processing component. The physical time rate of progression matches that of the underlying hardware and therefore allows an interested party to compare the rate of the physical time with the rate of its own processing time or media time progression in order to match processing rates. In concert with querying the current media time, and with feedback as to stream processing starvation or flooding, a processing component can not only make adjustments for matching processing rates, but also determine if it is ahead or behind a desired stream position for the master or reference stream (<i>i.e.</i> , synchronize).
Correlated Media Time	This property returns both the current media time value for the clock mechanism, and a corresponding reference time value as a single atomic operation. In the exemplary embodiment, the PC time serves as a reference time that is universally available to many different operational entities. By using the correlated media time, a processing component using a different clock may perform a translation using the common PC time while introducing a very small, and in many instances insignificant, amount of error. Since translations back and forth would use the same atomic operation, the small error is not cumulative.
Correlated Physical Time	This property returns both the current physical time value for the clock mechanism, and a corresponding reference time value as a single atomic operation. In the exemplary embodiment, the PC time serves as a reference time that is universally available to many different operational entities. By using the correlated physical time, a processing component using a different clock may perform a translation using the common PC time while introducing a very small, and in many instances insignificant, amount of error. Since translations back and forth would use the same atomic operation, the small error is not cumulative.
Granularity & Error	This property returns the clock increment granularity or resolution in order to indicate how precise or fine each clock "tick" will appear. In the exemplary embodiment, the granularity value is the number of 100 nanosecond increments per each clock tick. If a clock runs at or under 100 nanosecond resolution or granularity, the value would simply return 1. The granularity property allows a client of the clock to react differently depending on the resolution. The property further provides an error indication in 100 nanosecond increments with 0 representing the least amount of error.
Parent	This property returns the unique identifier of the processing component which produced the clock mechanism. This allows a client or processing component to determine if the clock mechanism is its own clock mechanism or some other processing component's clock mechanism so that it may react differently as the circumstances dictate. For the exemplary embodiment, this property would be a handle to the file object of an input connection pin instance or filter upon which the clock mechanism was produced.
Component State	This property reflects the state of the underlying processing component (<i>i.e.</i> , filter) so that a clock mechanism client may determine if the underlying processing component is in the stop state, the pause state, the run state, or the data acquisition state. Again, this property functions similar to the media time in that it is dependent upon what is happening to the underlying data stream.

TABLE 4 (continued)

Clock Mechanism Properties	
Properties	Description
Function Table	This property indicates a data structure or table that allows a client to access a table of functions which represent a subset of the property set interface. Of necessity, this property is platform specific since kernel mode function calls cross filter boundaries may be handled differently by different operating systems. In the exemplary embodiment, using the NT operating system, the function table allows other kernel mode entities to have a faster interface available at the DPC level. This is contrasted with the normal method of using properties through NT kernel IRPs.

Those skilled in the art will appreciate that other properties than those listed in Table 4 or a subset of the properties listed in Table 4 may be used to make a clock mechanism property set that provides functionality according to the interconnected filter system explained herein. Furthermore, the flexible property set mechanism allows a filter developer with unique timing capabilities to extend a "required" property set by adding thereto another property set, having its own GUID, for more advanced capabilities. The set mechanism allows the flexibility to filter developers in that a minimum amount of functionality may be implemented so that the filter is compliant with the system architecture while allowing unbounded ability for the filter developer to customize and add additional capabilities.

While a property set allows other processing components, such as kernel mode filters, to query relevant time properties, clock mechanisms may also support an event set so that notifications may be sent directly to interested clock mechanism clients by way of event notification IRPs. The handle of a file object may be "registered" with the clock mechanism as part enabling an event so that the clock mechanism will know where to send event notifications. Optionally, direct function call handles may be used so that DPC level processing can occur between kernel mode entities for achieving higher performance.

Table 5 below, a number of possible notification events make up an event set that can be supported by compliant drivers. A third party controlling agent may interconnect or register those entities for receiving events based on a given filter graph topology.

TABLE 5

Clock Mechanism Events	
Event	Description
Interval Event	An event which is to be signalled in a cyclic manner. In the exemplary embodiment, a specific interval and specified clock time are set by a controlling agent. Beginning at the specified clock time, cyclic event notifications will be sent at the occurrence of each interval. For example, if a processing component renders video frames from a video stream at thirty frames per second, it would specify a start time and an interval of 1/30th of a second so that the event notifications would be sent at the same frequency thereby allowing the video frames to be rendered properly.
Positional Event	An event which is to be signalled at a specified clock time. For example, a processing component or filter which renders a media event stream may have the need, in some circumstances, to play notes at a specific clock time. The processing component or filter could specify a positional event notification to occur at the designated clock time so that it would know to begin rendering the note. Positional notification events could also be used to determine exactly when the clock begins to progress after being put into the run state by setting an event notification on the stopped or paused clock to be signalled at the current time.

Those skilled in the art will appreciate that other timing related events may be created for a clock mechanism. For example, events may be broken down according to the different kinds of time (*i.e.*, media time or physical time) that a clock mechanism may support. Furthermore, "hybrid" events may be envisioned by combining elements of interval events and positional events, as well as other events based on state changes.

Referring now to Figure 16A, a method for synchronizing one or more "slave" processing components or filters based on a "master" media processing time is presented. After starting at step 380, a slave processing component or filter receives the master clock media time value at step 382. Such receipt may occur by having the slave filter or component query the master clock mechanism for the media time value or the master clock mechanism may send notifica-

tions to the slave processing component or filter. Applying the method explained by the flow chart of Figure 16A to the filter graph topology shown in Figures 15A and 15B, the master clock mechanism 344 is represented by a file object which would communicate with the input pin instance 336 of the video renderer filter 332 as a slave processing component.

Note also, that the media time value used depends upon the master data stream processing. Should media rendering be paused or stopped, then time would effectively stop for synchronization purposes.

The slave processing component or filter will compare the master clock media time value to the slave data stream media time value at step 384 in order to determine how far ahead or behind the slave media processing has become. Note that media time is based upon the time interval information in the actual stream of data samples and therefor can be viewed as positional in nature. In other words, synchronization also means processing two data streams at the same relative time position within the stream.

Before ending at step 388, the slave processing component or filter will adjust slave media processing to match the master media processing time at step 386. Such adjustments may include accelerating or decelerating slave media stream processing, changing the rate of the underlying hardware processor, etc.

Decelerating media stream processing rate may be achieved by duplicating samples for processing, introducing intermittent time delays, or resampling. Accelerating media stream processing may be achieved by omitting media samples for processing, or resampling. Those skilled in the art will appreciate that numerous ways exist for adjusting the media processing rate once it is determined that a change must be made.

Note that in the example topology shown in Figures 15A and 15B, input pin instance 336 would receive timing notifications for the video renderer filter 332. The video renderer filter 332 would make whatever adjustments to processing that are necessary in order to bring the processing of the video stream of Figure 15A in synchronization with the processing of the audio stream of Figure 15B.

Referring now to the flow chart of Figure 16B, processing steps for rate matching a live or real time audio stream origination rate with a processing component or filter processing rate is shown. After beginning at step 390, physical time samples are received at the processing component or filter at step 392. Again, this can be done by querying a clock mechanism or by receiving a timing event notification. With reference to Figure 15B, processing of the live audio stream 352 which is ultimately rendered at audio renderer filter 364, the relevant processing component would be the audio renderer filter 364 as event notifications and control information are received at the input pin instance 368. Also, the clock mechanism 344 would provide the different timing information.

From the physical time samples received at step 392, the amount of data processed between those time samples can be used to compute or determine a physical time rate at step 394. The physical time rate is the rate at which the audio hardware actually renders data and represents data rendering throughput and is also known as the filter processing rate. If the source of the live data stream produces data samples at a faster origination rate than the filter processing rate, excess data occurs. If the origination rate is less than the filter processing rate, then gaps or starvation for data processing occur. In either case, performance degradation results.

The actual processing rate of the processor that originates the data samples, or ordination rate, can be determined from the media or data samples themselves. Because the media samples have time interval information in the form of convention or actual timestamp information, media time samples may be computed from the media stream at step 396. This allows the origination rate representing the processing rate of the hardware creating the media samples to be computed by taking a fixed amount of data processed and dividing this quantity by time it took for that data to be processed as determined by the media time samples.

Before ending at step 400, the media stream processing is adjusted to match the ordination rate to the filter processing rate at step 398. Again, adjustments may occur by omitting samples for rendering, adding time delays between sample rendering, and other ways known in the art.

Referring now to the flow chart shown in Figure 16C, processing steps for using a clock mechanism to make a translation is shown. After beginning at step 402, a processing component or filter will receive a correlated time value from a clock mechanism at step 404. Note that the processing component or filter will have access to the reference time clock, such as a PC clock, that generates the reference time value portion of the correlated time value. A media time delta is computed at step 406 by subtracting the media time value from the PC clock or other reference time value as found from the correlated time value received previously. The media time delta will be used by the processing component or filter for translation and represents the variation of the media time from the PC clock time.

In order to make the translation, the processing component or filter gets a current PC clock time value (or other reference time value) at step 408. The current media time at the processing component or filter may be set or timestamp information may be translated using the current PC clock time value and the master media time delta at step 410. Since the media time delta shows the variation of the media time from PC clock time in the reference time basis, it may be simply added to the current PC clock time value in order to arrive at a proper "translated" or new time value for timestamp information. After finalizing the translation process in step 410, processing ends at step 412. If this communication were taking place between two components on a remote processor, such as a DSP, the reference time value would be

the local host time rather than the PC time.

Applying the method for translation illustrated in the flow chart of Figure 16C to the interconnected kernel mode filters for processing and rendering a live audio and video stream as shown in Figures 15A and 15B, it would be the input pin instance 336 on the video renderer filter 332 that would receive the correlated time from the clock mechanism 344 at step 404. Finally, it would be the video renderer filter 332 that made the computations and translation of either the timestamps on the live video stream 320 or otherwise created a translated reference time based on the media time of the clock mechanism 344.

Those skilled in the art will recognize that in a system of interconnected filters as disclosed herein, the uniform buffer allocation facilities and the timing facilities may be used together in combination for efficient processing or used separately depending implementation problem solution requirements. Furthermore, the timing facilities may be used independently of a system of interconnected filters.

Those skilled in the art will recognize that the methods of the present invention may be incorporated as computer instructions stored as a computer program code means on a computer readable medium such as a magnetic disk, CD-ROM, and other media common in the art or that may yet be developed. Furthermore, important data structures found in computer hardware memory may be created due to operation of such computer program code means.

Claims

1. A method for interconnecting software drivers to allow efficient kernel mode processing of data and having a standardized way of providing timing information for synchronizing different data streams or rate matching different hardware clocks, the method comprising the steps of:

opening one or more kernel mode drivers;

forming one or more connection pin instances for connecting the drivers, each connection pin instance hierarchically related to one of said one or more drivers and used for data transmission between said one or more drivers;

creating one or more clock mechanisms for rate matching and stream synchronization, each clock hierarchically related to one of said one or more connection pin instances and providing a data stream time and a physical time based on an underlying hardware oscillator; and

interconnecting said one or more connection pin instances so as to provide a continuous data flow path through the one or more drivers residing in kernel mode, said clocks making timing information available for rate matching between different clock oscillations and for synchronizing different data streams.

2. A method as claimed in claim 1, in which each connection pin instance is represented by a file object and the hierarchical relationship is created by specifying the related driver, the driver referenced as a file object of an I/O device available on the system, during connection pin instance file object creation and each clock mechanism is represented by a file object and the hierarchical relationship with a connection pin instance is created by specifying the related connection pin instance file object as a parent during the clock file object creation.

3. A method as claimed in claim 1, which includes the step of querying each of said connection pin instances on each of said one or more drivers to determine the availability of clock mechanisms in order to determine whether to instantiate a clock mechanism for synchronization or rate matching on a particular connection pin instance prior to interconnecting said one or more connection pin instances.

4. A method as claimed in claim 1, in which at least one of said one or more connection pin instances on said one or more drivers support at least one predefined set of properties, set of methods, and set of events to indicate to a third party component the availability of a clock mechanism on the at least one connection pin instance and for allowing a third party component to form said clock mechanism on respective connection pin instances and the clocks support at least one predefined set of properties, set of methods, and set of events to control said respective clocks as determined by the third party component.

5. A method for synchronizing multiple streams of data comprising the steps of:

creating a master clock reference based on a reference data stream of data samples, each data sample having time interval information indicating position within the reference data stream, the master clock reference created from the time interval information; and

processing one or more other data streams based on the master clock reference so that the other data streams are synchronized with the processing of the reference data stream.

6. A method as claimed in claim 5, in which the master clock reference is accessed as part of a property set.
7. A method as claimed in claim 5, in which the master clock reference is part of a selectively instantiable clock mechanism that is created in response to a request by a third party agent.
8. A method as claimed in claim 5, in which the time interval information is provided by timestamp information on the data samples.
9. A method as claimed in claim 5, in which the time interval information is provided media stream convention.
10. A method for rate matching a data stream processed through a processing component to a hardware processor comprising the steps of:
 - providing, to the processing component, a physical time reference based on the hardware oscillations of a hardware device; and
 - adjusting the rate of processing at the processing component of a data stream to match the hardware device processing rate based on the physical time reference, the data stream comprised of samples having time interval information associated therewith.
11. A method as claimed in claim 10, in which the physical time reference is accessed as part of a property set.
12. A method as claimed in claim 10, in which the physical time reference is part of a selectively instantiable clock mechanism that is created in response to a request by a third party agent.
13. A method as claimed in claim 10, in which the time interval information is provided by timestamp information on the data samples.
14. A method as claimed in claim 10, in which the time interval information is provided data stream convention.
15. A method as claimed in any one of claims 10 to 14, in which the hardware device is a hardware processor.
16. A method as claimed in any one of claims 10 to 14, in which the hardware device is a hardware renderer which renders the samples in the data stream.
17. A method for translating a designated time value in a first component based on a designated time value in a second component and utilizing a common hardware oscillator comprising the steps of:
 - receiving, at a first component from a second component, a correlated time value in a single operation, the correlated time value comprising a designated time value and a common time value based a common hardware oscillator;
 - querying, by the first component, the current value of the common hardware oscillator; and
 - translating, by the first component, the designated time value in the first component based on the current value of the common hardware oscillator and the correlated time value received from the second component.
18. A method as claimed in claim 17, in which the designated time value is a positional time value based on the position within a data stream having time interval information associated therewith.
19. A method as claimed in claim 17, in which the designated time value is a physical time value based upon a hardware oscillator.
20. A kernel mode data processing system allowing synchronous processing of two data streams comprising:
 - a first data source;
 - a first plurality of kernel mode data processing components including an originating component and a terminating component, the originating component reading data samples of a first data stream originating from the first data source, and at least one of the processing components having a master clock component associated therewith;
 - a second data source;

a second plurality of kernel mode data processing components including an originating component and a terminating component, the originating component reading data samples of a second data stream originating from the second data source, and at least one of the processing components synchronizing processing based on the master clock component associated with a processing component found in the first plurality of processing components; and

kernel mode component interconnections between the data processing components of respective first and second plurality of kernel mode data processing components to route the processing of respective data samples from the respective originating component to the respective terminating component.

21. A data processing system as claimed in claim 20, in which:

each of said first plurality of kernel mode data processing components and each of said second plurality of kernel mode data processing components comprise at least one connection pin instance;

said kernel mode interconnections between said first plurality of data processing components are formed by interconnecting the at least one connection pin instance of one of said first plurality of data processing components to another of said first plurality of data processing components; and

said kernel mode interconnections between said second plurality of data processing components are formed by interconnecting the at least one connection pin instance of one of said second plurality of data processing components to the at least one connection pin instance of another of said second plurality of data processing components.

22. A data processing system as claimed in claim 21 or claim 22, in which the terminating components are renderers for rendering data.

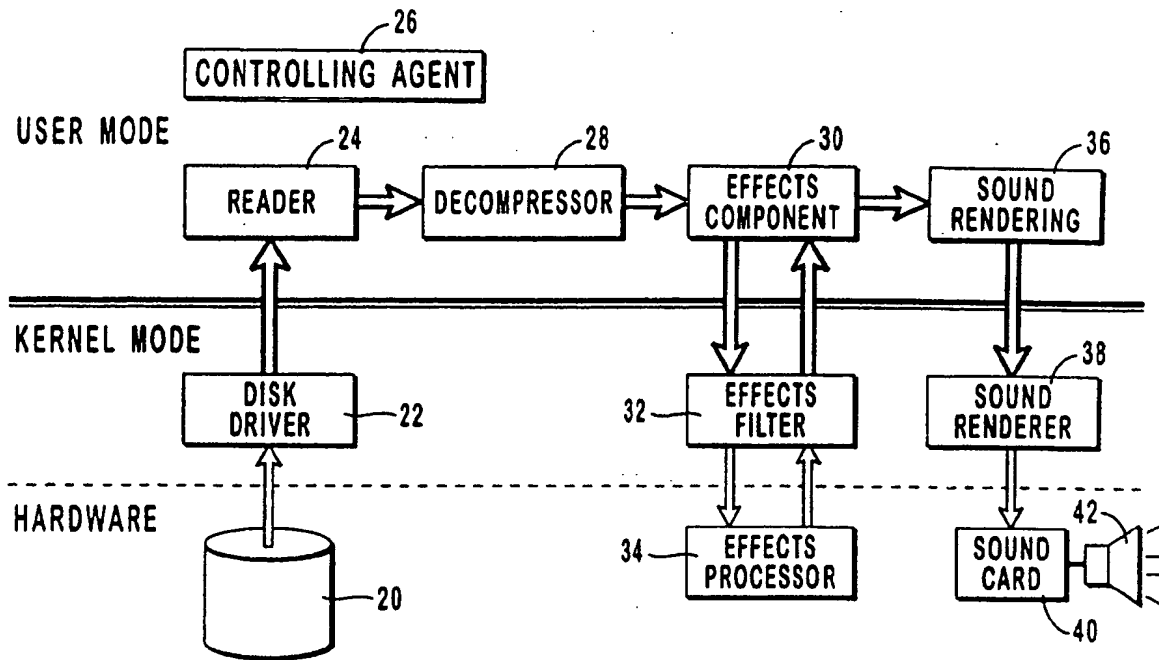


FIG. 1
(PRIOR ART)

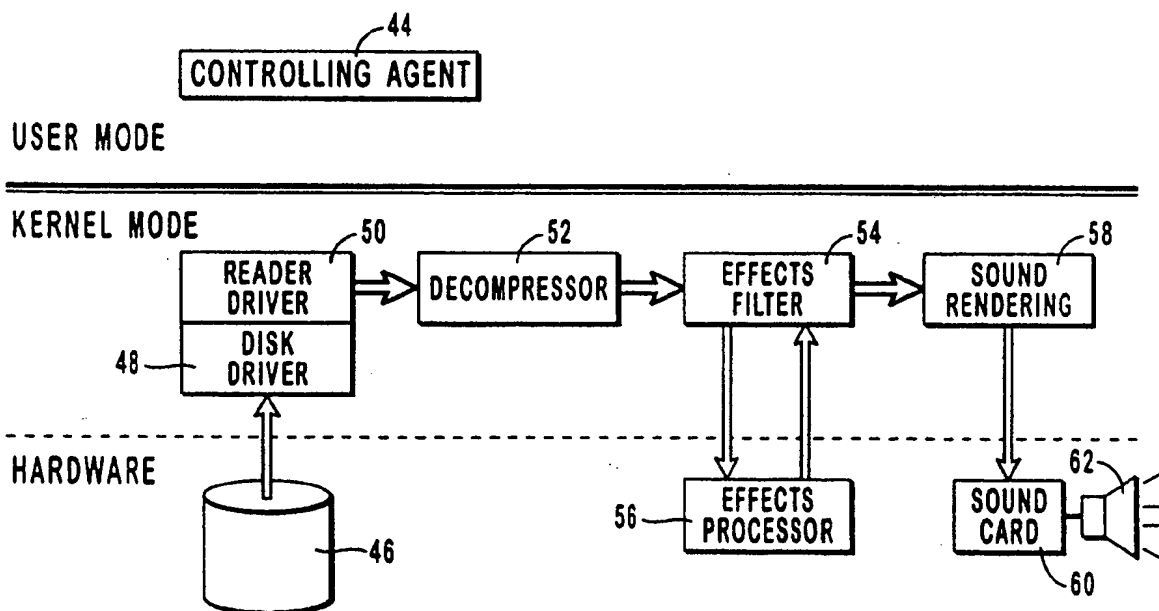


FIG. 2

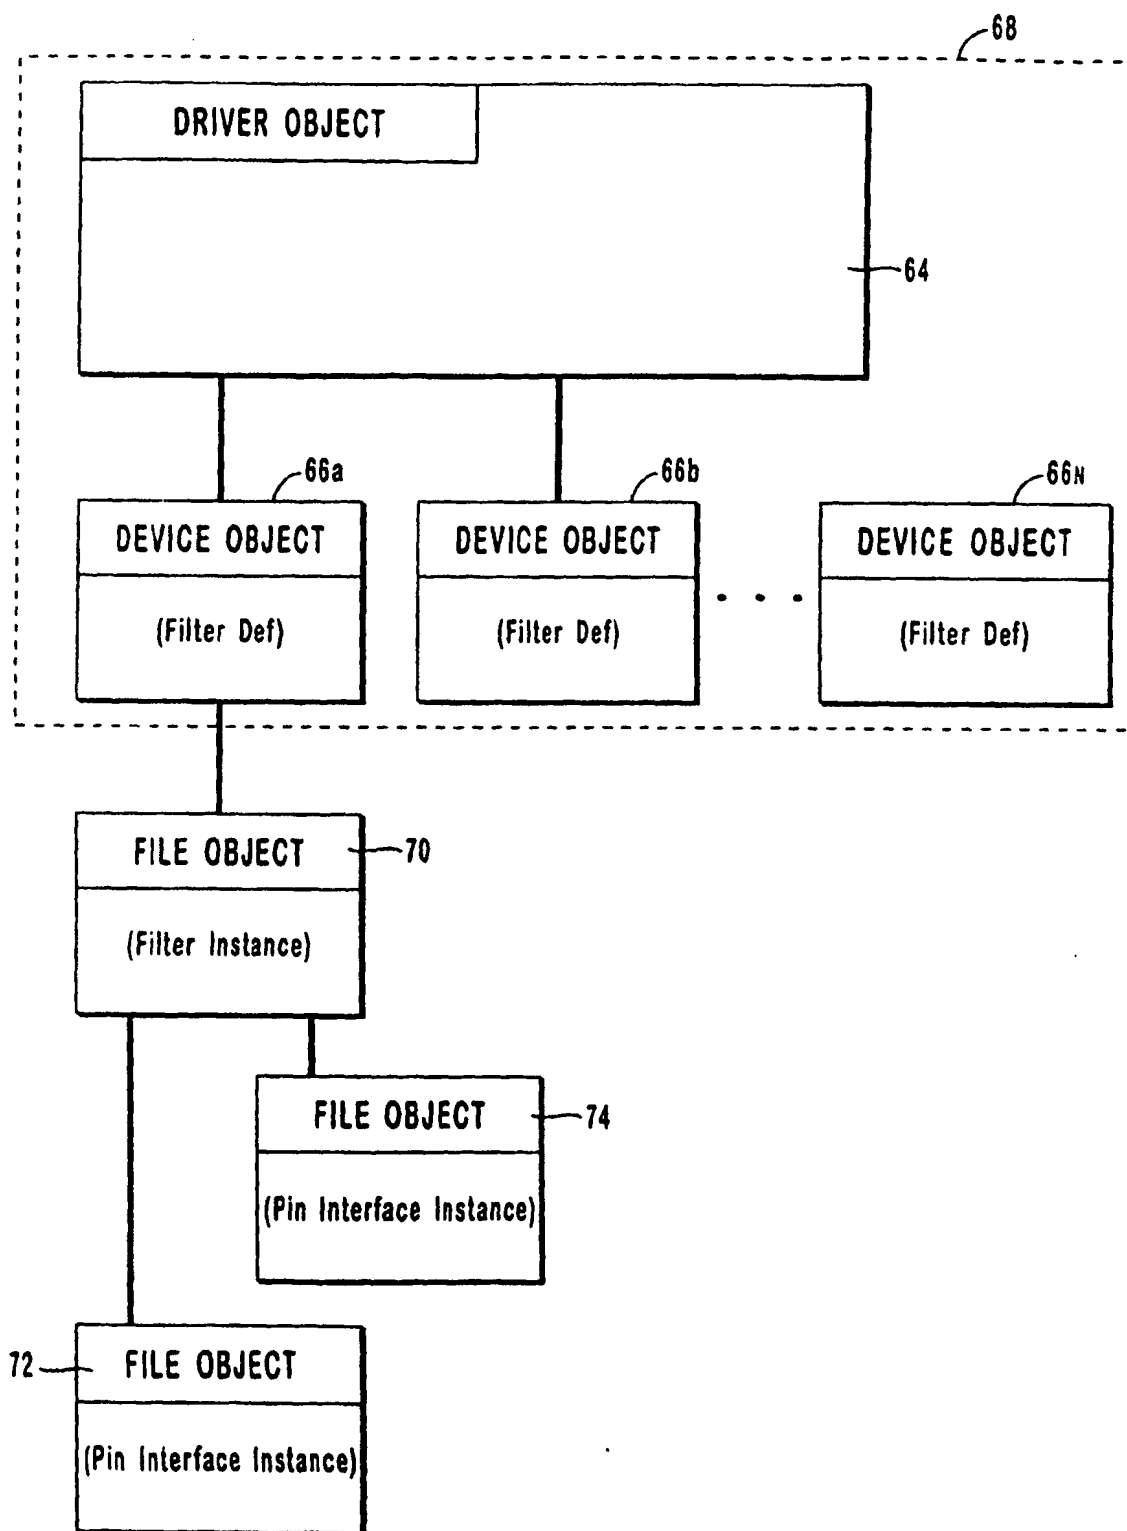


FIG. 3

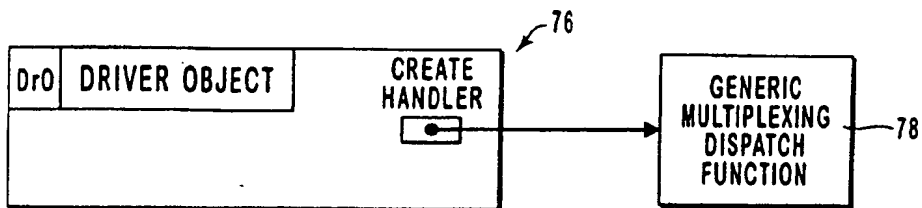


FIG. 4A

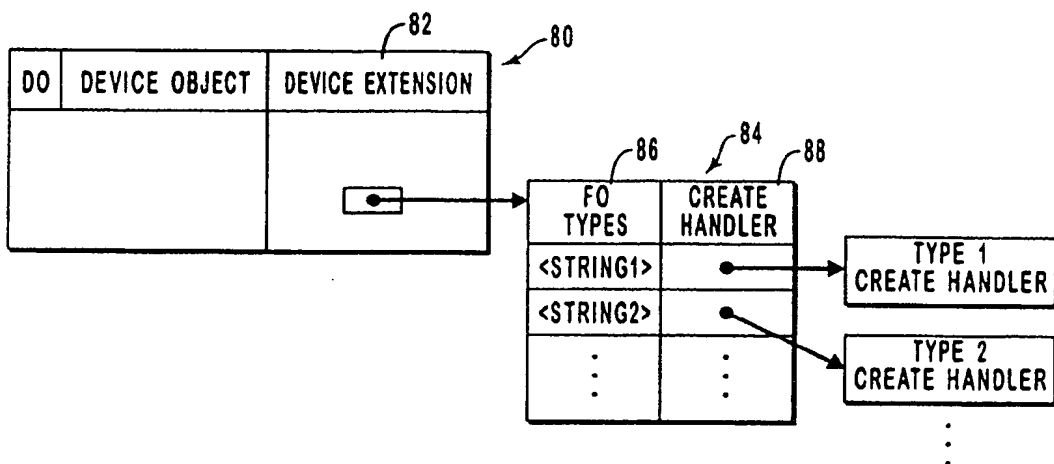


FIG. 4B

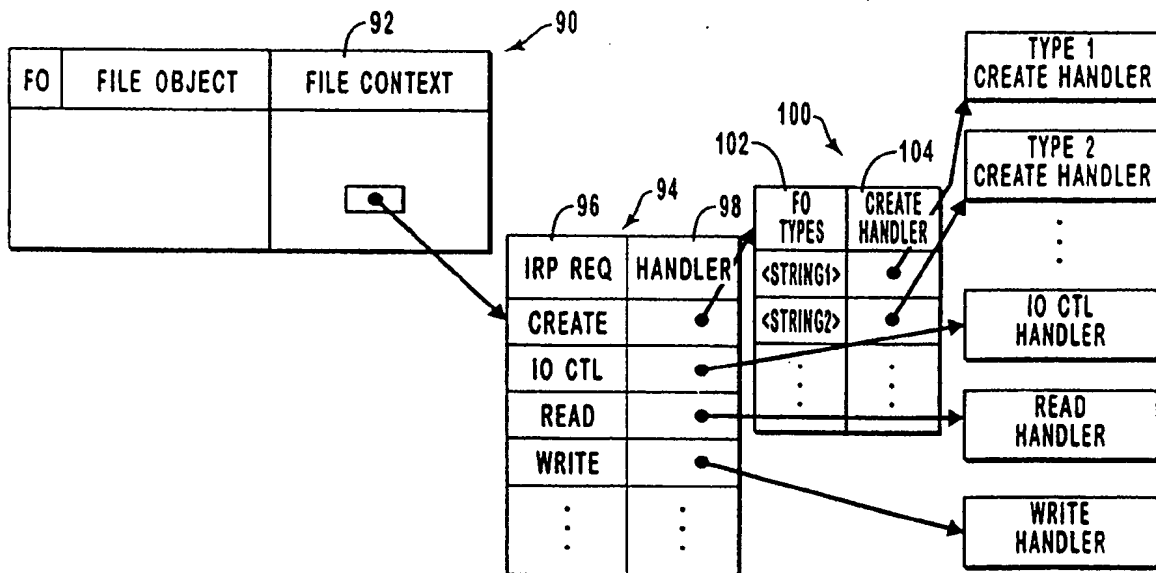


FIG. 4C

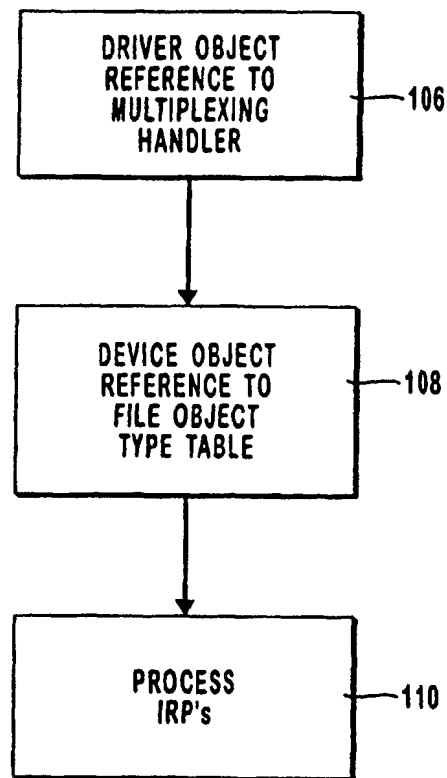
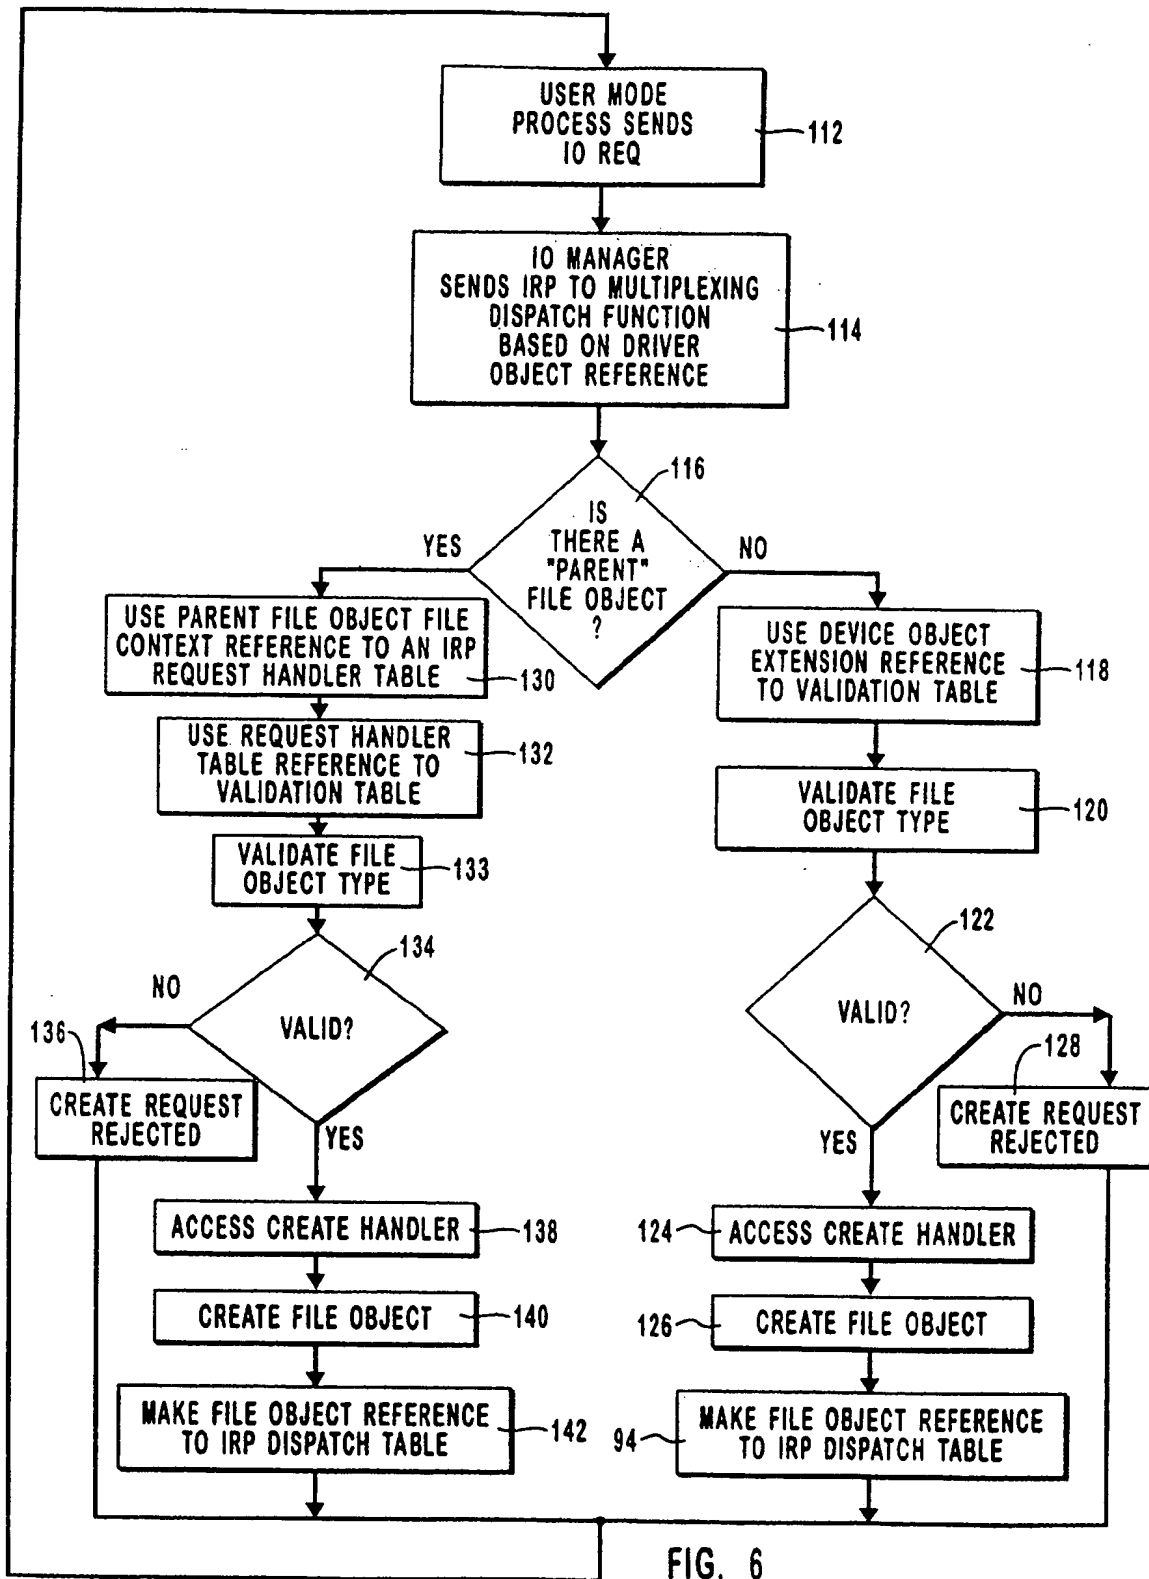


FIG. 5



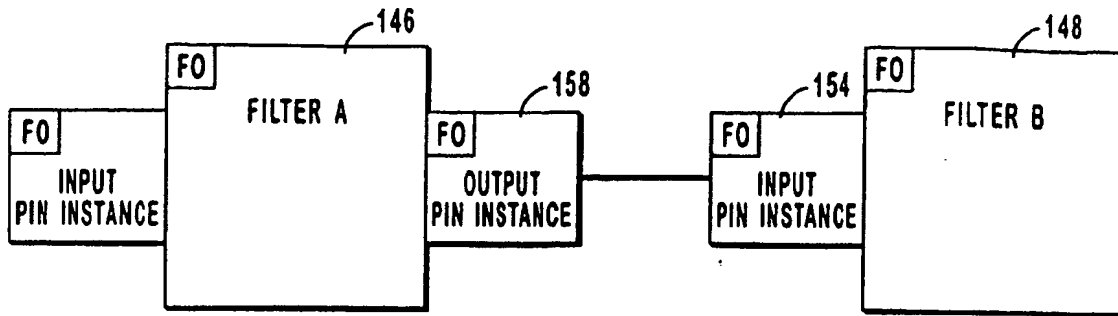


FIG. 7

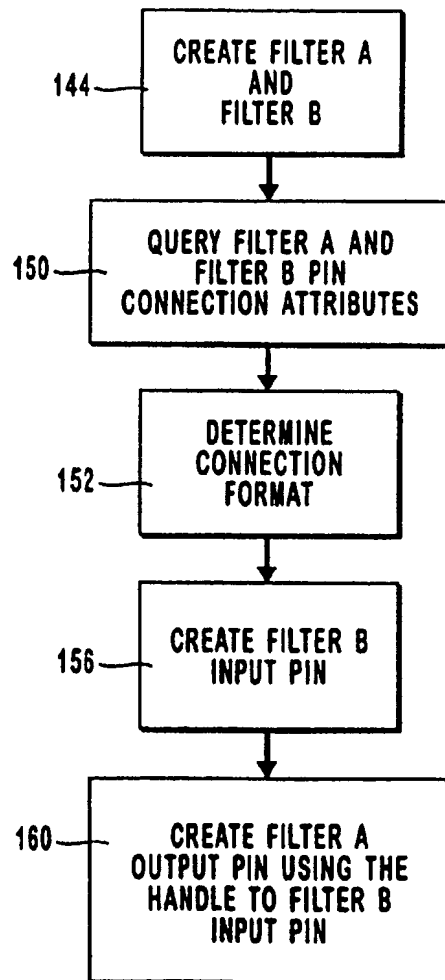


FIG. 8

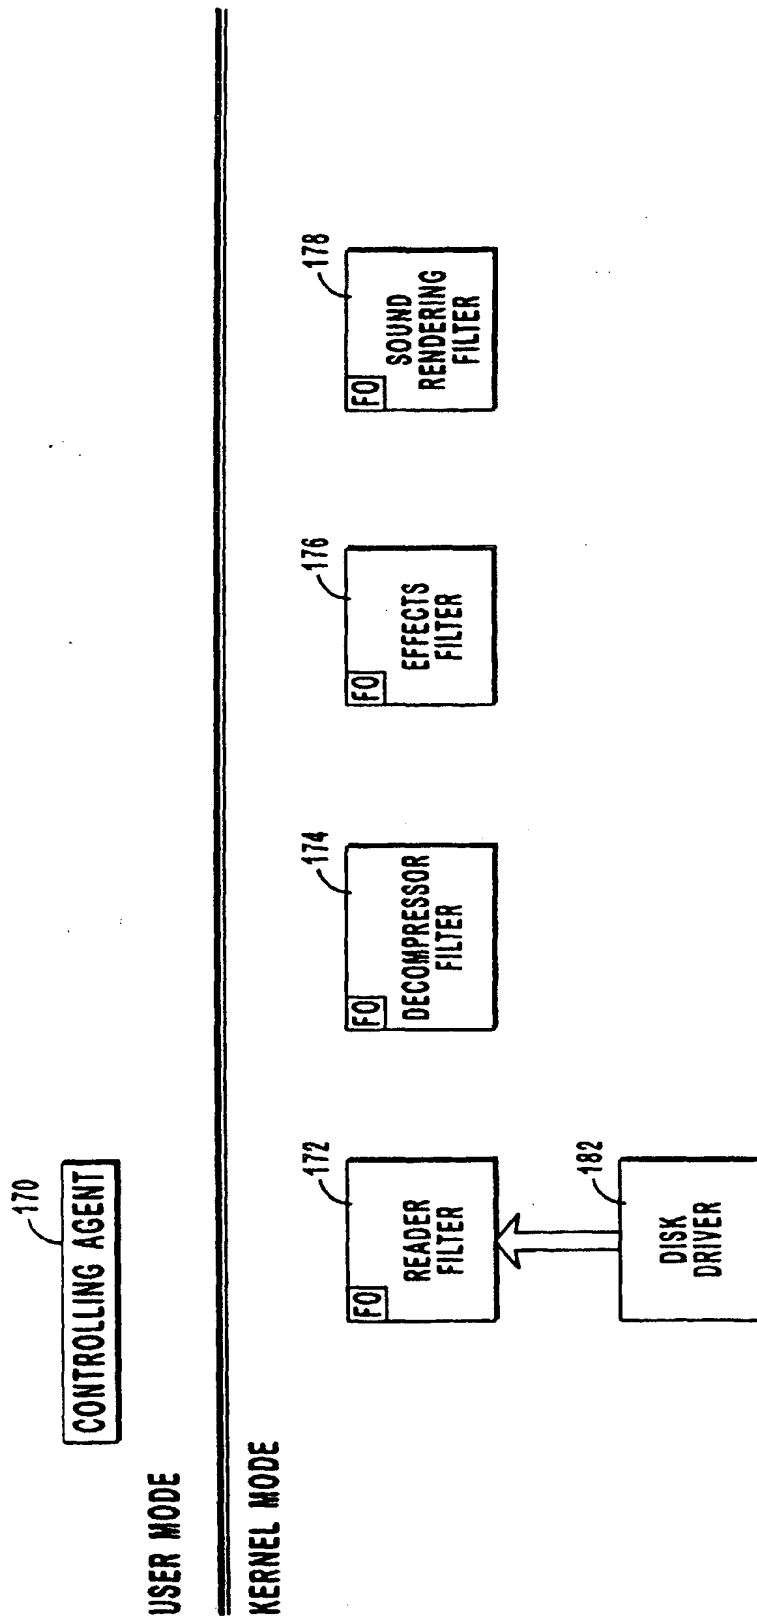


FIG. 9A

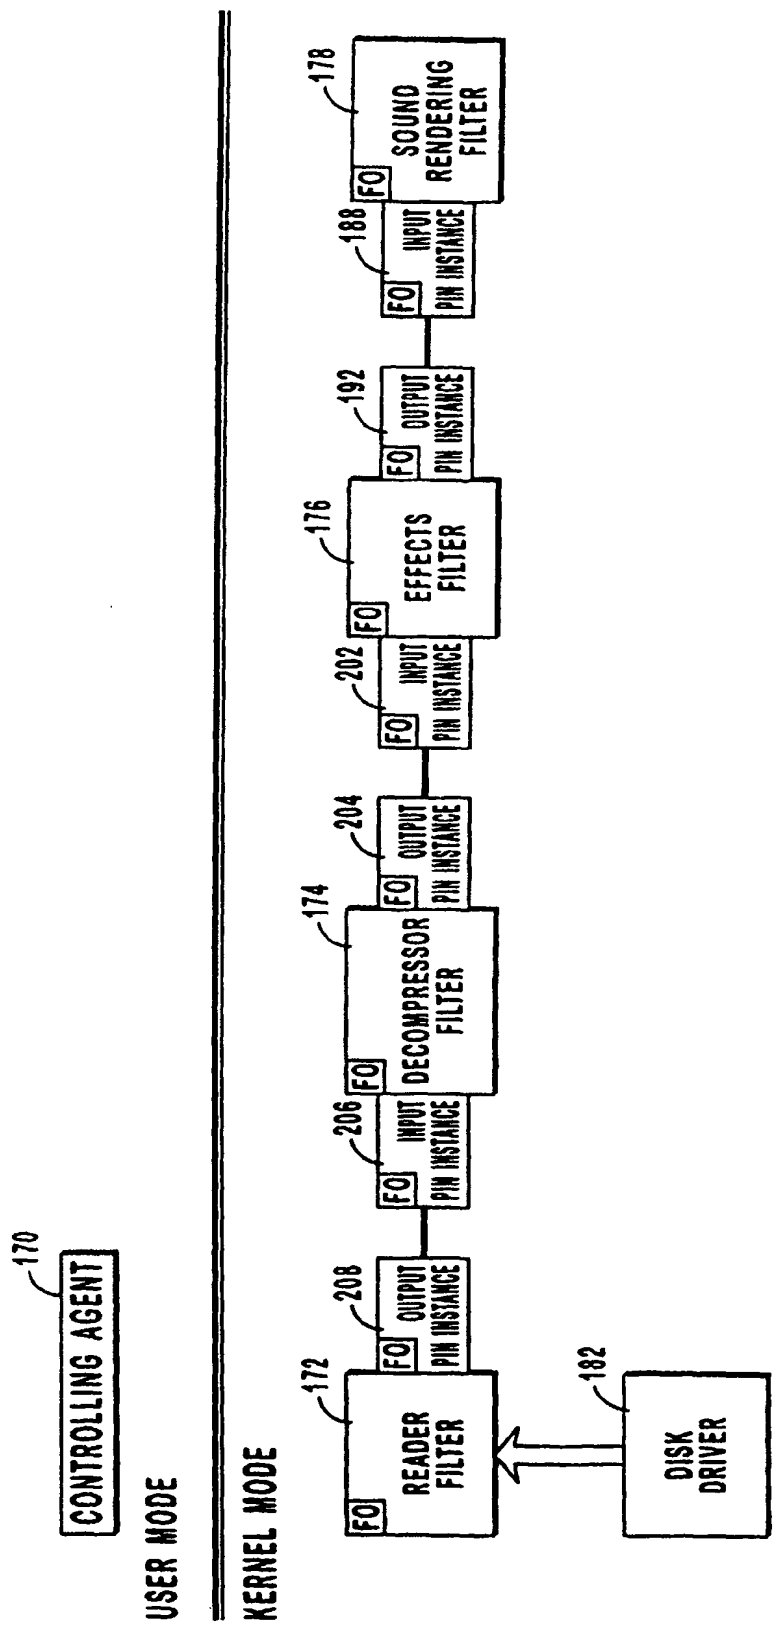


FIG. 9B

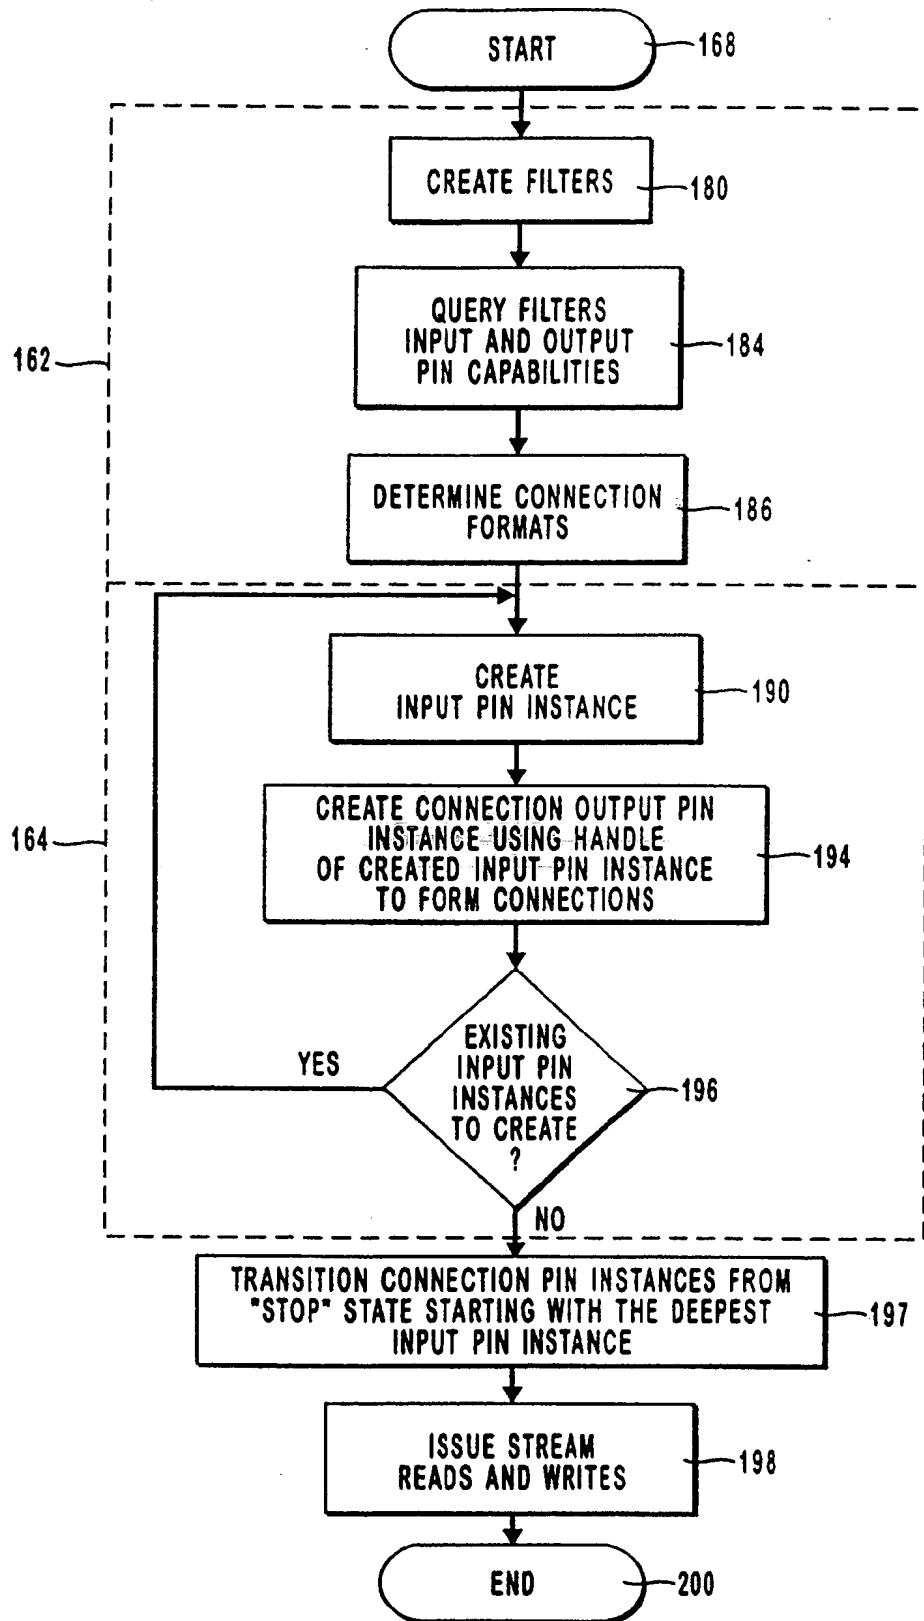


FIG. 10

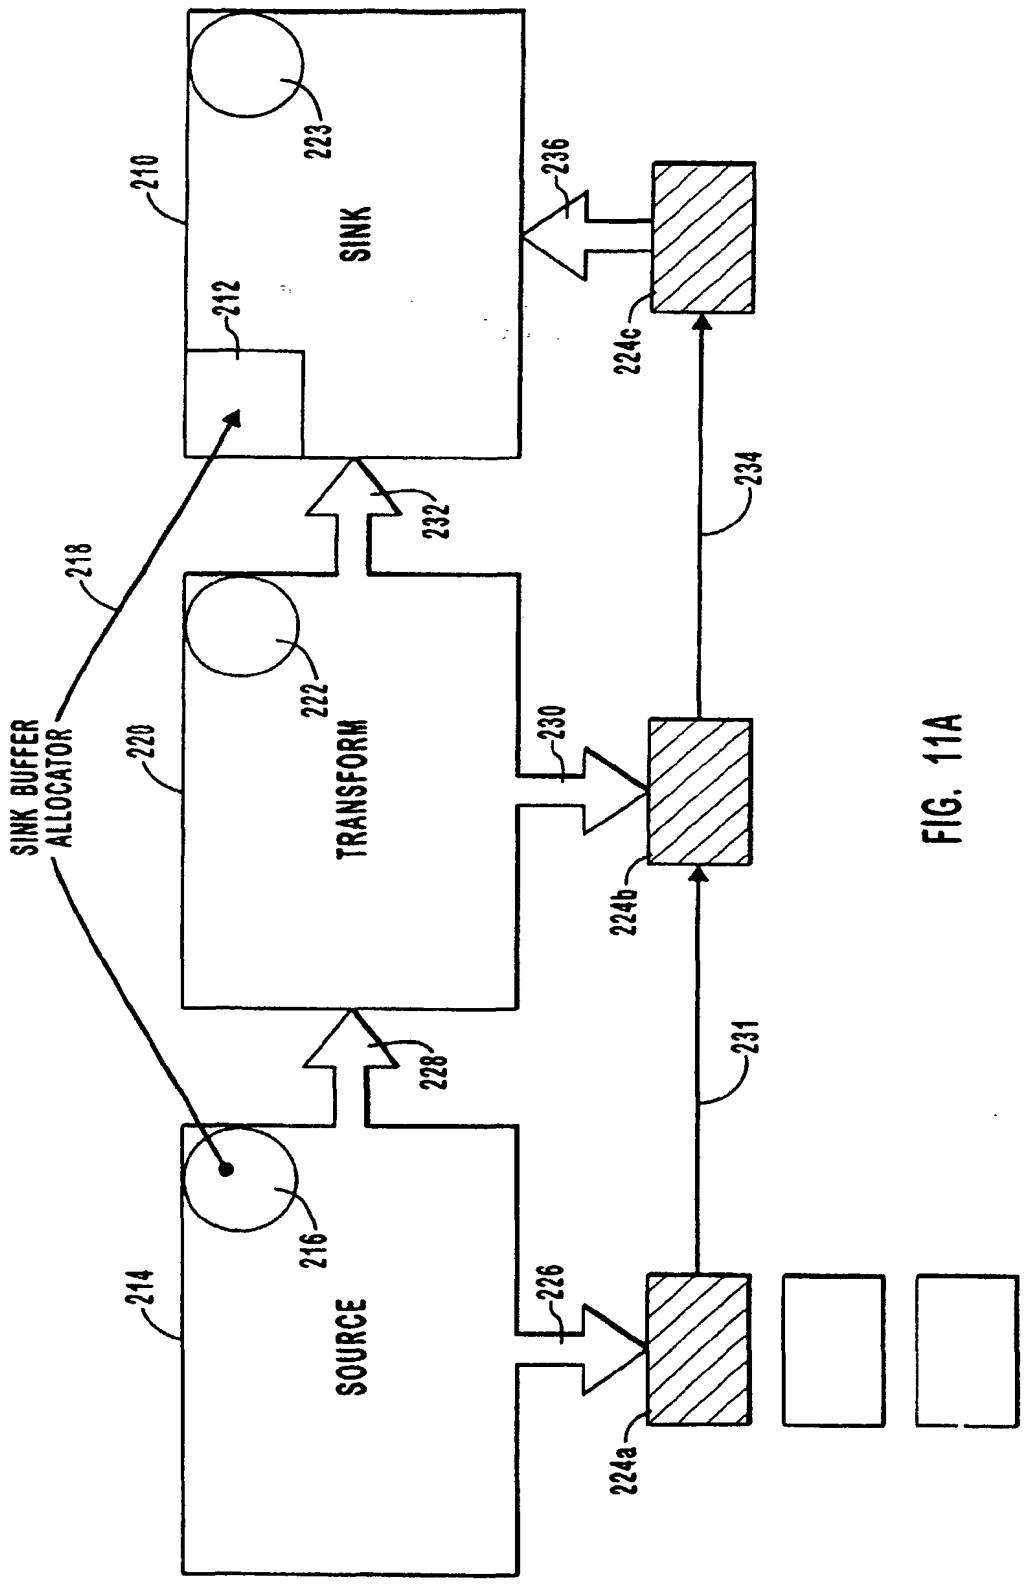


FIG. 11A

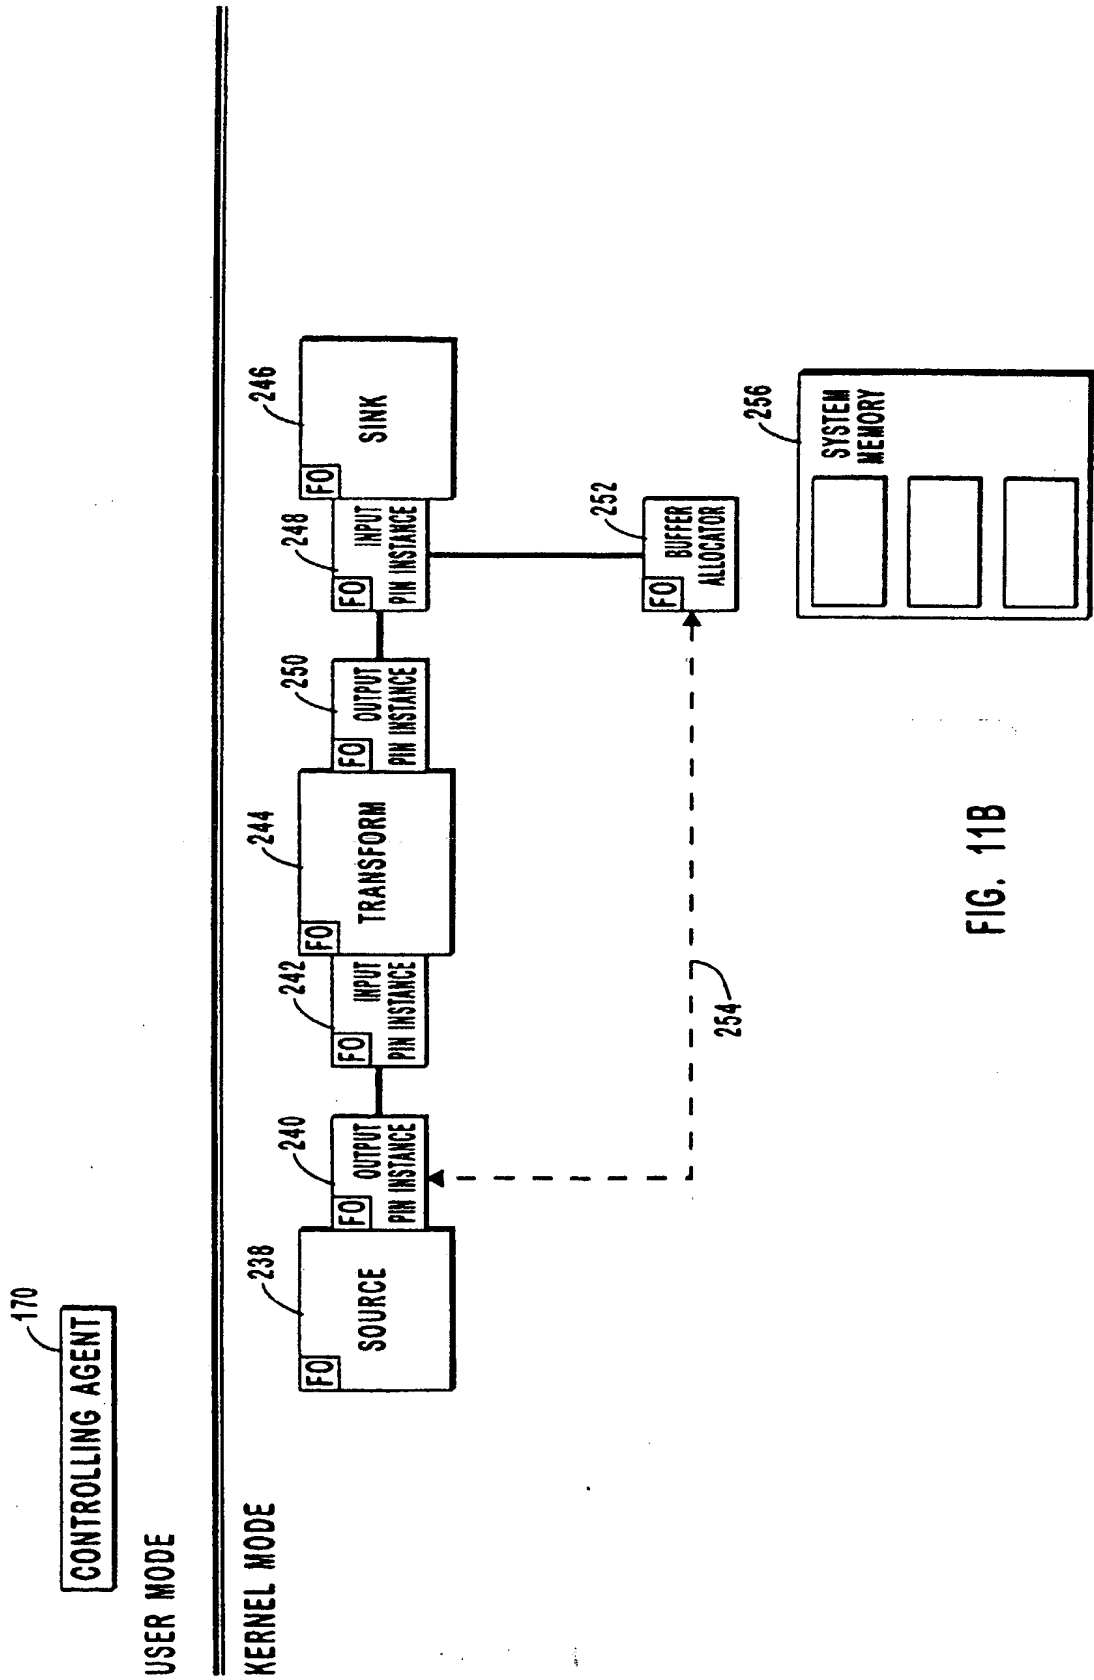


FIG. 11B

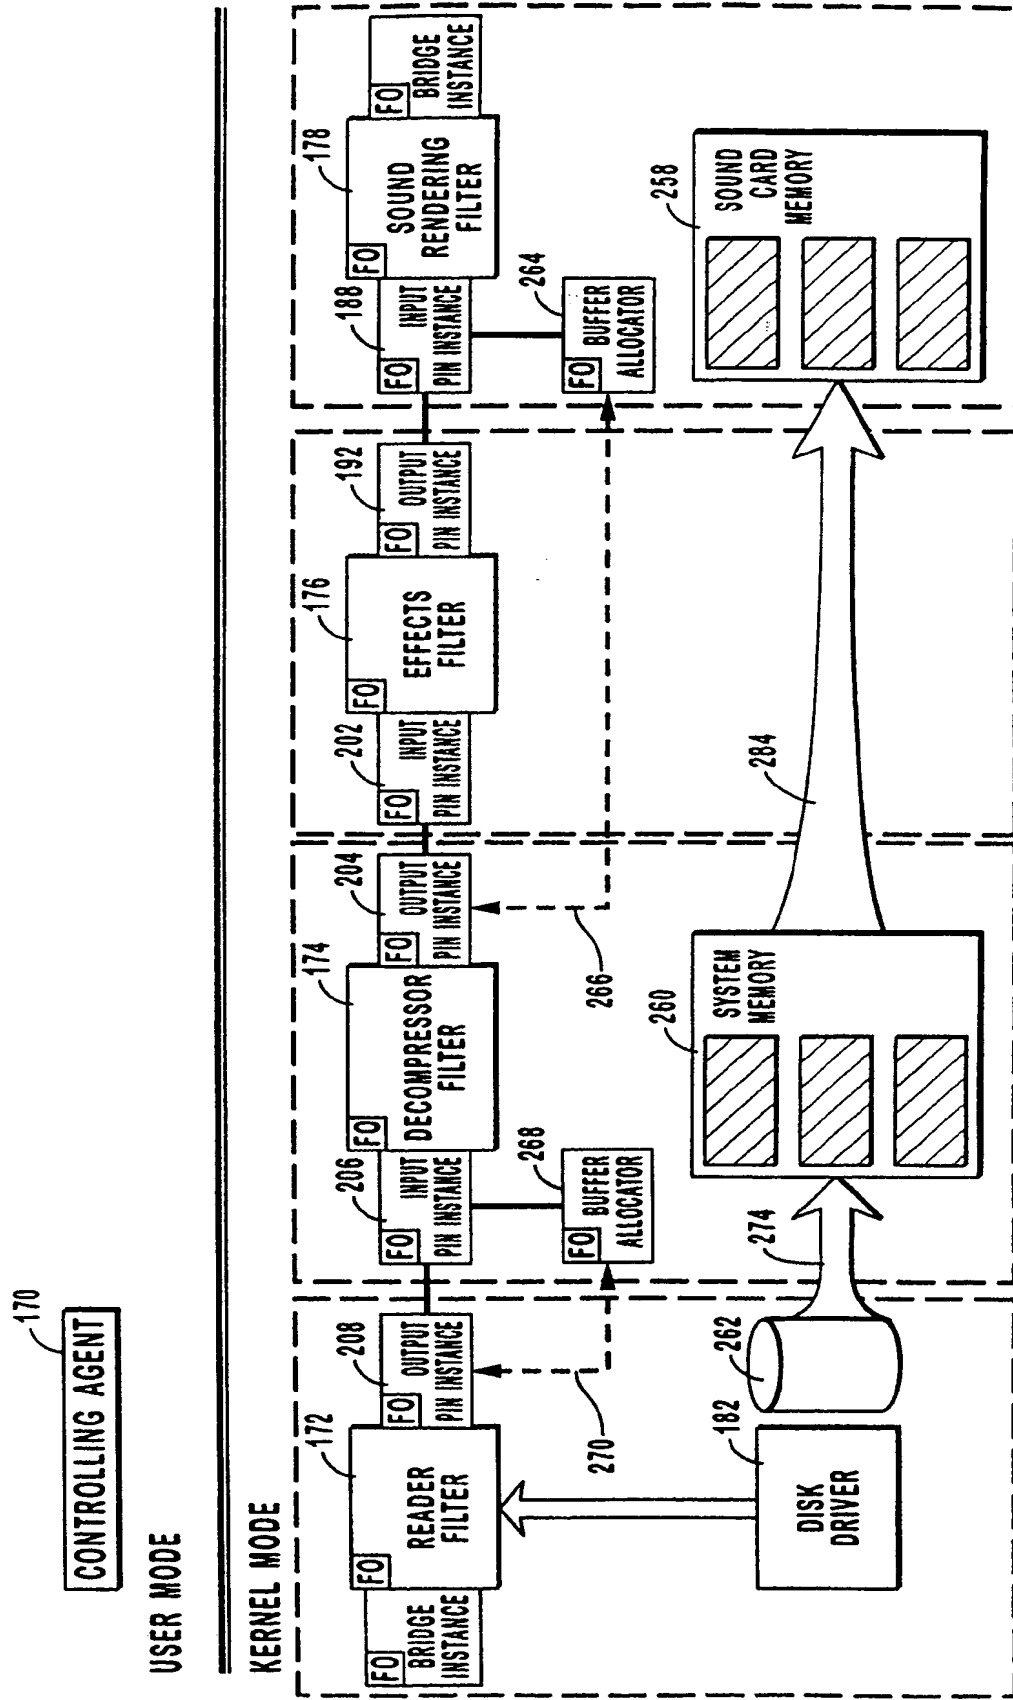


FIG. 12

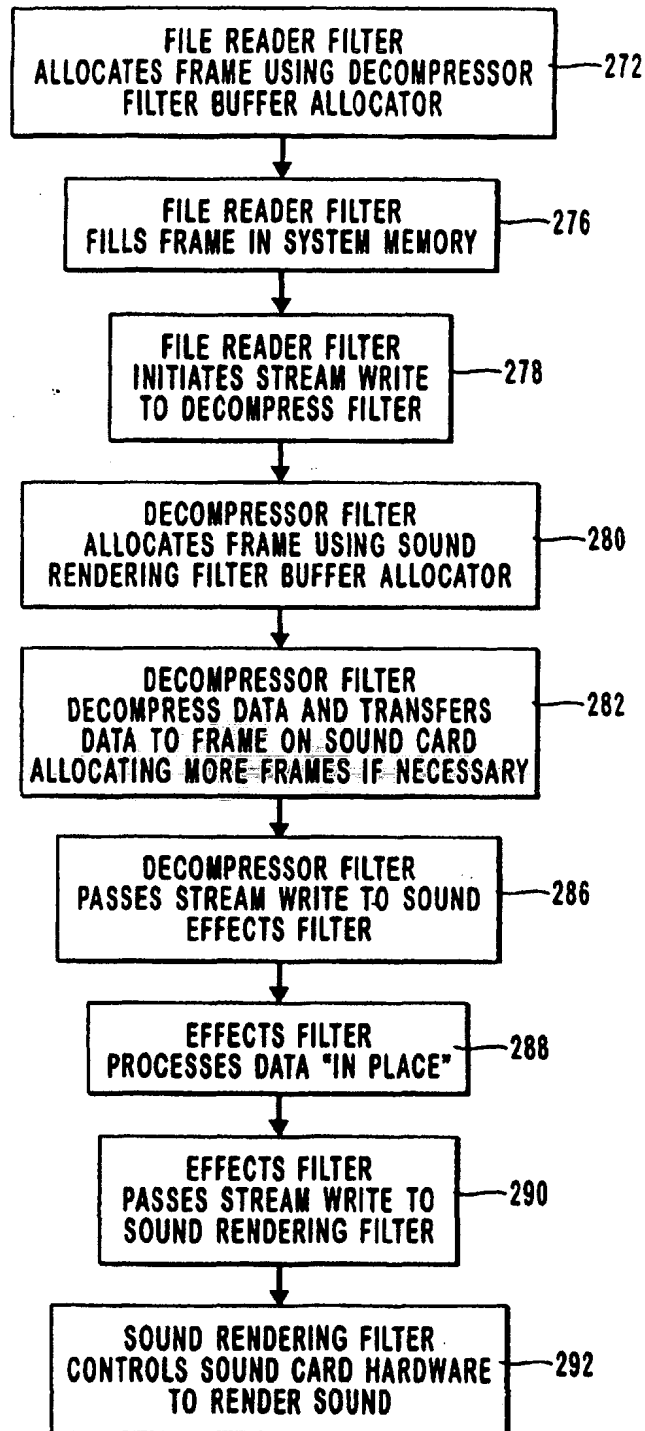


FIG. 13

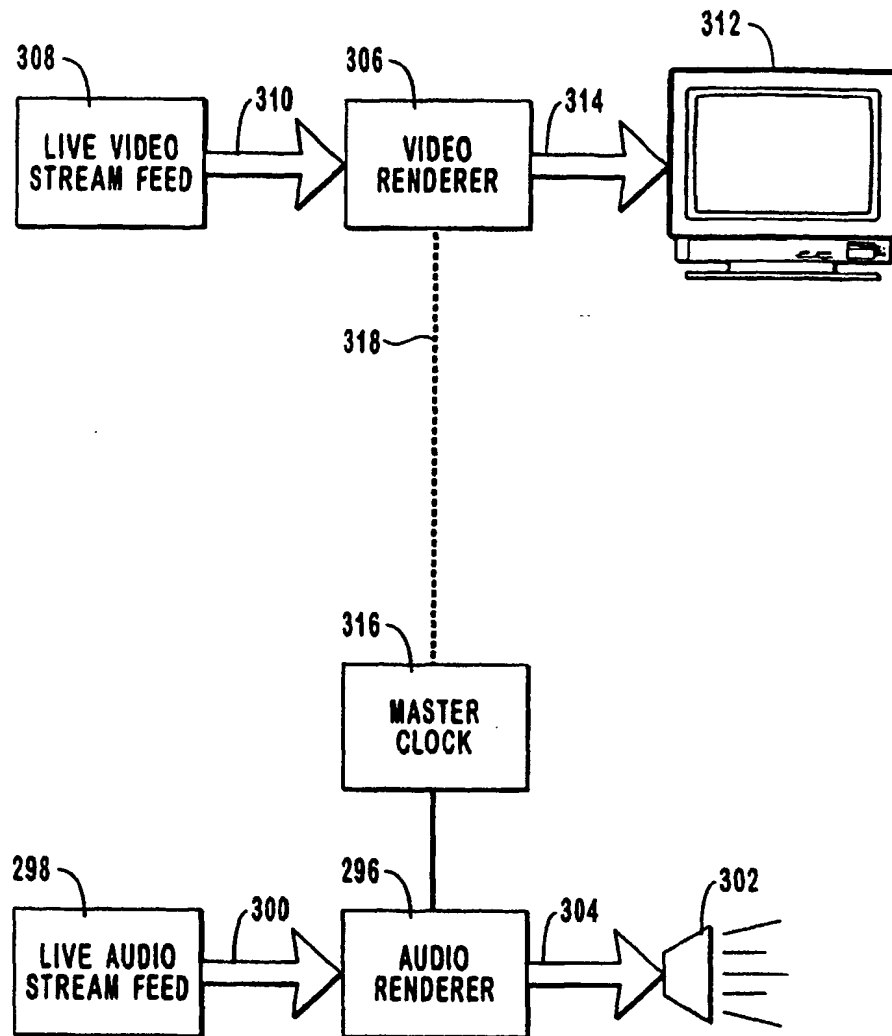


FIG. 14

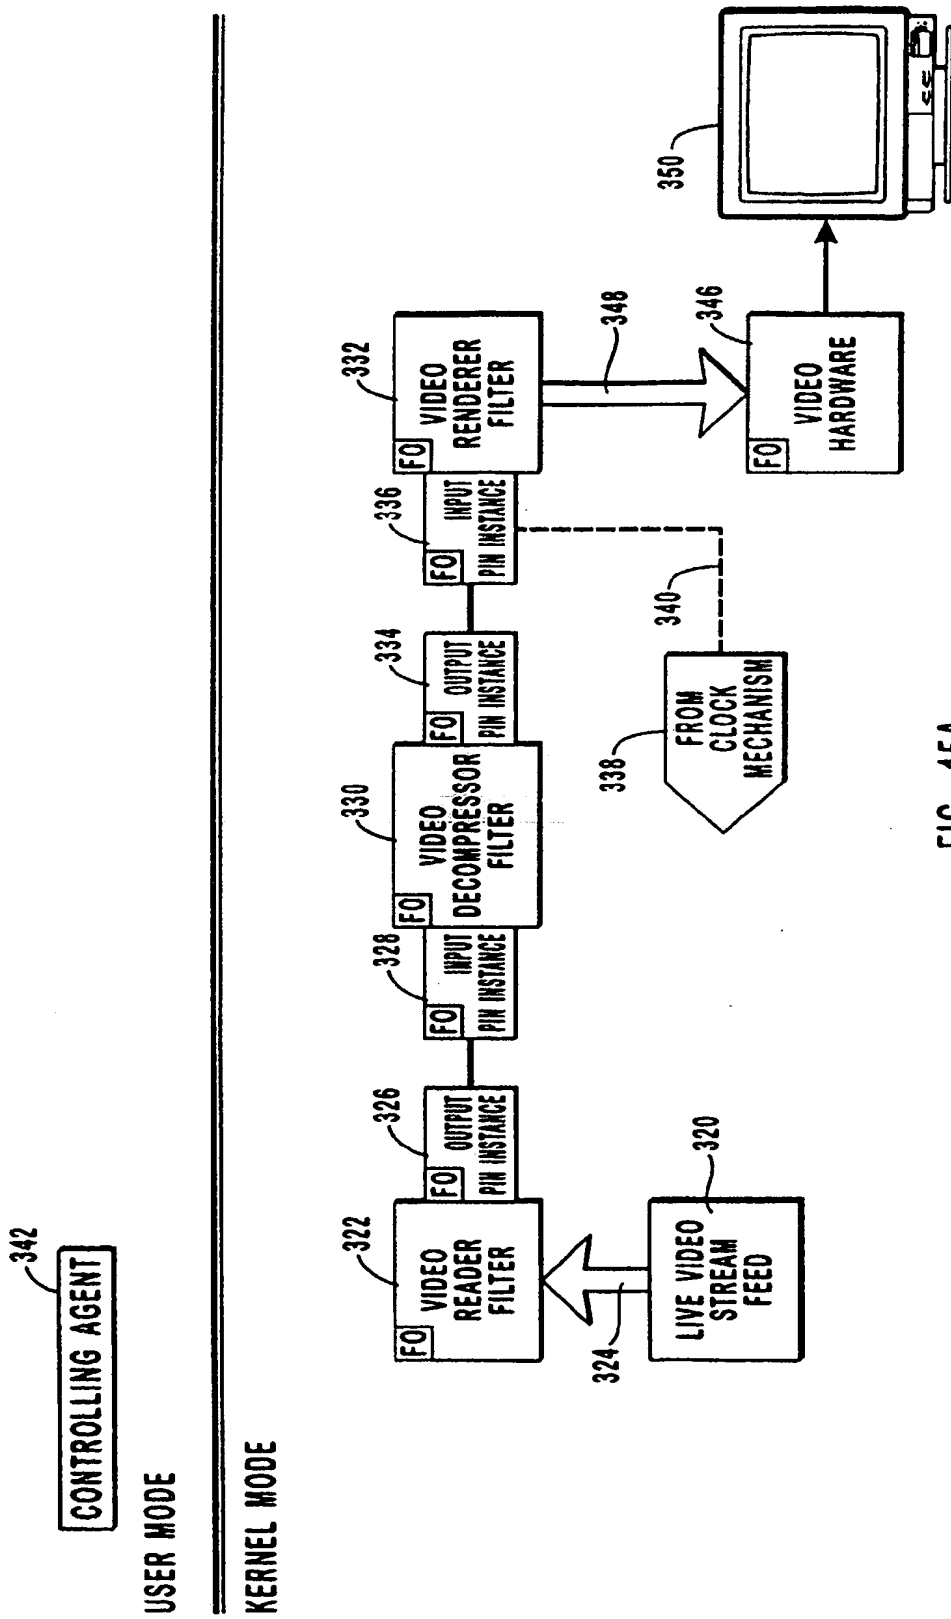


FIG. 15A

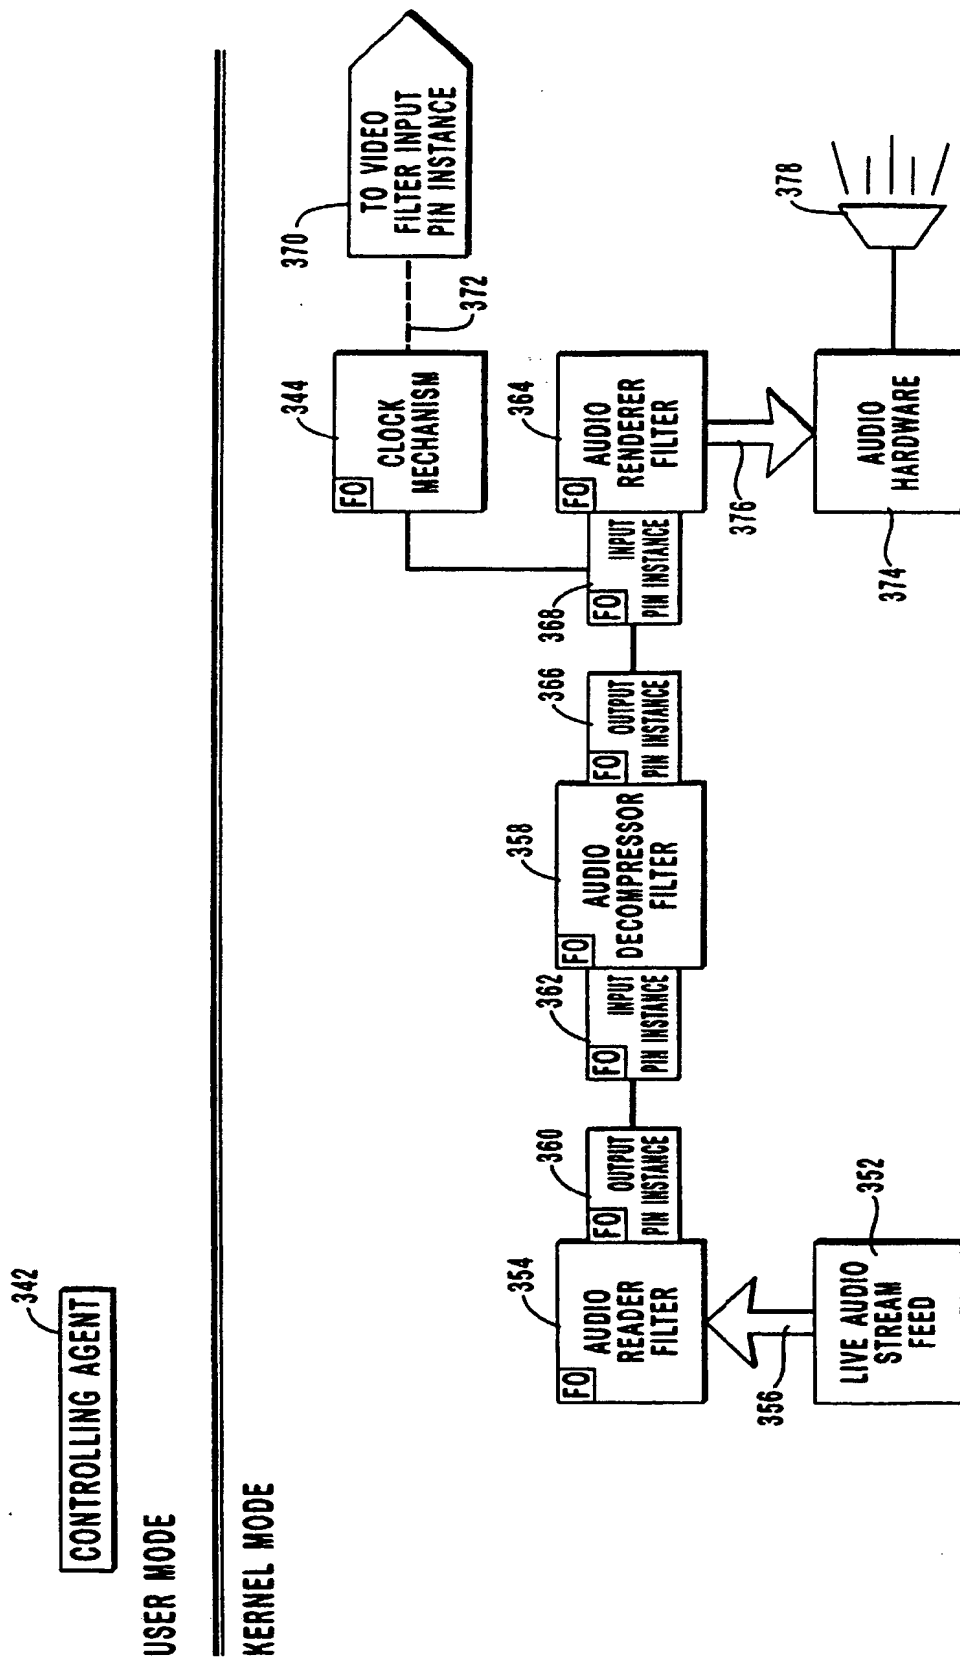


FIG. 15B

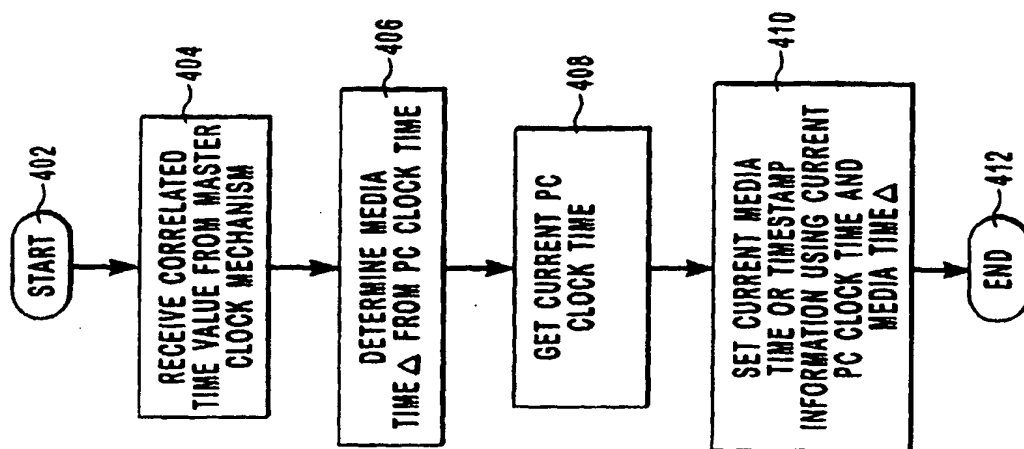


FIG. 16C

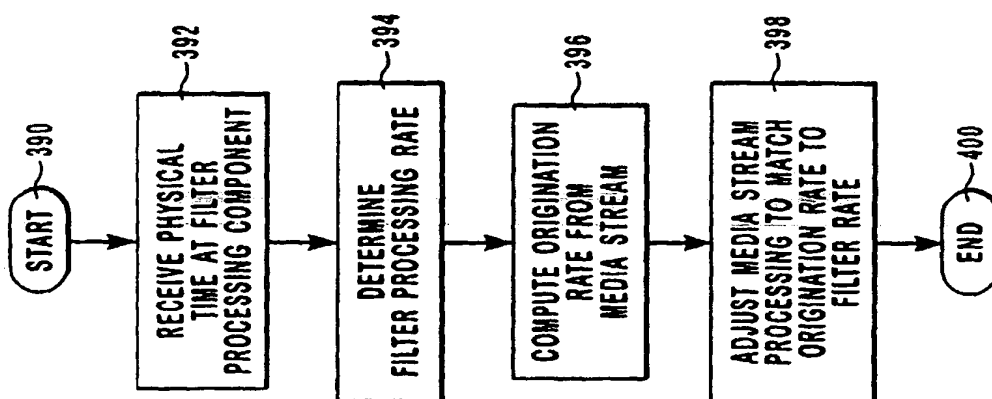


FIG. 16B

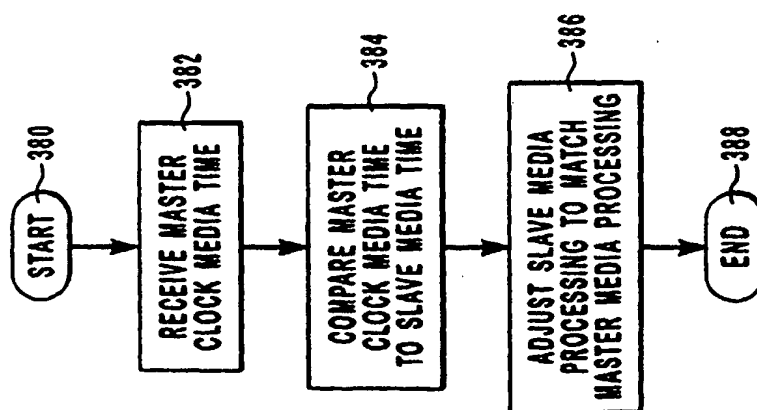
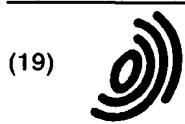


FIG. 16A

THIS PAGE BLANK (USPTO)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 0 869 443 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
09.05.2001 Bulletin 2001/19

(51) Int. Cl.⁷: **G06F 17/30**

(43) Date of publication A2:
07.10.1998 Bulletin 1998/41

(21) Application number: **97304294.8**

(22) Date of filing: **19.06.1997**

(84) Designated Contracting States:
**AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE**
Designated Extension States:
AL LT LV RO SI

(30) Priority: **04.04.1997 US 826560**

(71) Applicant: **MICROSOFT CORPORATION**
Redmond, Washington 98052-6399 (US)

(72) Inventors:
• **Shaw, George H.J.**
Woodinville, WA 98072 (US)

• **O'Rourke, Thomas J.**
90800 Oulu (FI)
• **Woodruff, Bryan A.**
New Bend, WA 98045 (US)

(74) Representative:
Belcher, Simon James et al
Urquhart-Dykes & Lord
Tower House
Merrion Way
Leeds LS2 8PA (GB)

(54) **Method and computer program product for synchronizing the processing of multiple data streams and matching disparate processing rates using a standardized clock mechanism**

(57) A method and computer program product for synchronizing processing between two or more data streams (e.g., video and sound input) and for rate matching between two different hardware clocks that may drift with respect to one another (e.g., an originating clock represented in a timestamped data stream versus a clock actually rendering the data) in a system of interconnected software drivers running in kernel mode. The present invention overcomes the coordination complexity and inaccuracies in the prior art by providing a clocking mechanism in a system wherein multiple drivers having input and output connection pin instances are chained together. The clocking mechanism synchronizes between data streams by providing a master clock on an input pin instance of a driver that is used to synchronize with other input pin instances on other drivers and "slave" clocks. Synchronization is achieved through event notification or stream position queries so that corresponding frames of data in separate streams are rendered together (e.g., video frames with corresponding sound track). Rate matching is achieved through monitoring a physical clock progression in comparison with a series of data stream timestamps thereby allowing adjustments to match the different clock rates. A common physical clock (e.g., PC clock) can be used as a reference for a component to

translate a particular clock time to a time shared by all components with a minimum of error.

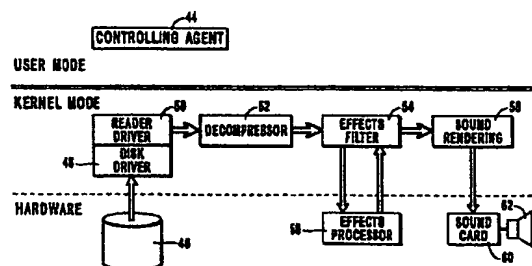


FIG. 2

EP 0 869 443 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 97 30 4294

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	ROTHERMEL K ET AL: "CLOCK HIERARCHIES: AN ABSTRACTION FOR GROUPING AND CONTROLLING MEDIA STREAMS" IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS,US,IEEE INC. NEW YORK, vol. 14, no. 1, 1996, pages 174-184, XP000548820 ISSN: 0733-8716 * the whole document *	1-5,8, 10,13, 15-21	G06F17/30
A	US 5 452 435 A (NEUMANN ERIK R ET AL) 19 September 1995 (1995-09-19) * column 4, line 38 - column 5, line 26 * * column 6, line 53 - column 12, line 50 * * claims *	1-5,10, 20,21	
A	HORN F ET AL: "ON PROGRAMMING AND SUPPORTING MULTIMEDIA OBJECT SYNCHRONIZATION" COMPUTER JOURNAL,GB,OXFORD UNIVERSITY PRESS, SURREY, vol. 36, no. 1, 1993, pages 4-18, XP000360261 ISSN: 0010-4620 * the whole document *	1,5,10, 17	TECHNICAL FIELDS SEARCHED (Int.Cl.6) G06F G11B
A	US 5 471 576 A (YEE RAYMOND L) 28 November 1995 (1995-11-28) * column 1, line 54 - column 2, line 19 * * claims *	5,10,17	
A	EP 0 598 499 A (IBM) 25 May 1994 (1994-05-25) * abstract * * figures *	1,5,10, 17	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 15 March 2001	Examiner Abbing, R
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 03/82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 97 30 4294

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

15-03-2001

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5452435 A	19-09-1995	AU 680488 B	31-07-1997
		AU 6418994 A	24-10-1994
		CA 2159466 A	13-10-1994
		DE 69425054 D	03-08-2000
		DE 69425054 T	08-03-2001
		EP 0692123 A	17-01-1996
		JP 8511891 T	10-12-1996
		WO 9423387 A	13-10-1994
US 5471576 A	28-11-1995	NONE	
EP 0598499 A	25-05-1994	US 5420801 A	30-05-1995
		JP 2522902 B	07-08-1996
		JP 7234795 A	05-09-1995

EPO FORM P0458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)